Double-Take® AVAILABILITY™

Version 7.1.1
Double-Take Availability for vSphere User's Guide

VISION® SOLUTIONS

Notices

Double-Take Availability for vSphere User's Guide Version 7.1.1, Thursday, January 29, 2015

- **Product Updates**—Check your service agreement to determine which updates and new releases you may be eligible for. Product updates can be obtained from the support web site at http://www.VisionSolutions.com/SupportCentral.
- **Sales**—If you need maintenance renewal, an upgrade license key, or other sales assistance, contact your reseller/distributor or a Vision Solutions sales representative. Contact information is available on the Vision Solutions Worldwide Locations and Contacts web page at http://www.VisionSolutions.com/Company/Vision-HA-Locations.aspx.
- **Technical Support**—If you need technical assistance, you can contact CustomerCare. All basic configurations outlined in the online documentation will be supported through CustomerCare. Your technical support center is dependent on the reseller or distributor you purchased your product from and is identified on your service agreement. If you do not have access to this agreement, contact CustomerCare and they will direct you to the correct service provider. To contact CustomerCare, you will need your serial number and license key. Contact information is available on the Vision Solutions CustomerCare web page at http://www.VisionSolutions.com/Support/Support-Overview.aspx .
- **Professional Services**—Assistance and support for advanced configurations may be referred to a Pre-Sales Systems Engineer or to Professional Services. For more information, see the Windows and Linux tab on the Vision Solutions Consulting Services web page at http://www.VisionSolutions.com/Services/Consulting-Services.aspx .
- **Training**—Classroom and computer-based training are available. For more information, see the Double-Take Product Training web page at http://www.VisionSolutions.com/Services/DT-Education.aspx.

Contents

Chapter 1 Agentless vSphere overview **4**
Replication process 5

Chapter 2 Agentless vSphere requirements **7**

Chapter 3 Creating an agentless vSphere job **12**

Chapter 4 Managing and controlling agentless vSphere jobs **30**
Viewing agentless vSphere job details 37
Validating an agentless vSphere job 40
Editing an agentless vSphere job 41
Viewing an agentless vSphere job log 42
Updating to a new controller appliance 44
Failing over agentless vSphere jobs 45

Chapter 5 Double-Take Console **46**
Double-Take Console requirements 48
Console options 49

Chapter 6 Managing servers **51**
Adding controller appliances 57
Adding replication appliances 58
Providing server credentials 60
Viewing server details 61
Editing server properties 63
Controller appliance associations properties 64
Server licensing 65
E-mail notification configuration 67
Viewing server logs 69

Chapter 1 Agentless vSphere overview

The fundamental function of Double-Take is to identify what you want to protect on your production server, known as the source, and replicate that to a backup server, known as the target. With agentless vSphere protection, there are multiple components which make up the source and target.

1. You need your hosts. (The use of the word host is used generically to represent a standalone ESXi host or a cluster, which is typically defined as a group of ESXi hosts.) One is your source host and one is your target host. In some cases, the source and target could be the same host.
2. You need virtual machines on your source host. These are the virtual machines you want to protect and are called your source virtual machines. Once you have established protection, you will have replicas of the source virtual machines on your target host. These are called your target virtual machines.
3. Each source and target host (or cluster) must have a replication appliance, which controls sending data between the hosts. These replication appliances are referred to as the source replication appliance and the target replication appliance.

Beyond your source and target components, you need a controller appliance. The controller appliance handles communication from the Double-Take Console to the replication appliances. It also handles licensing for your agentless vSphere jobs. The controller appliance should not be located with your source host, in case your source host experiences a failure. Ideally, the controller appliance should be located with your target host or on its own.

You must also have the Double-Take Console in your environment.

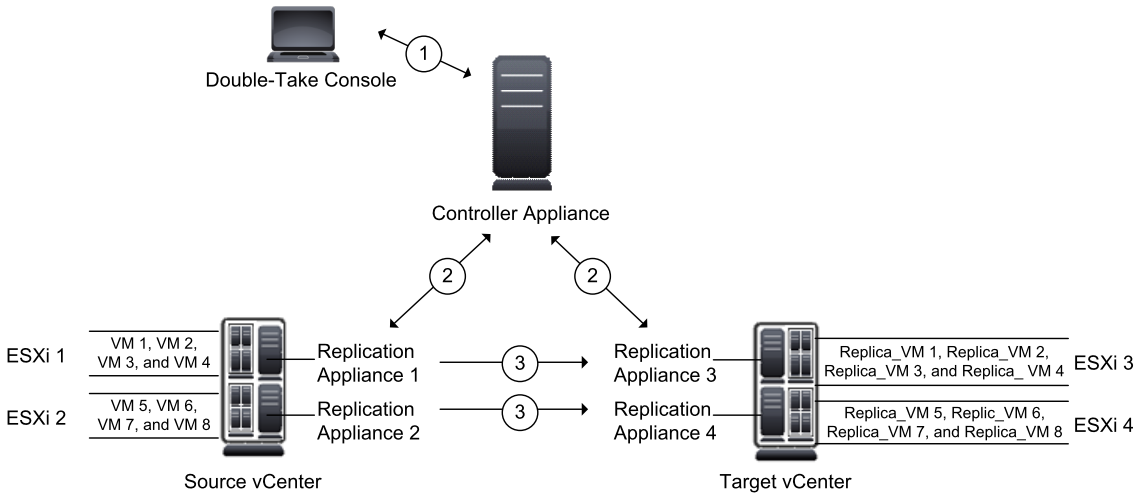

1. The Double-Take Console communicates with the Double-Take controller appliance.
2. The controller appliance communicates with the various Double-Take replication appliances.
3. The replication appliances on the source take the replication snapshots and send them to the target. The replication appliances on the target apply the replication snapshots and take point-in-time snapshots of the target. For more details on the replication process, see *Replication process* on page 5.

Replication process

Agentless vSphere protection is different than other Double-Take job types because it is reliant on ESXi technology. Instead of real-time replication, agentless vSphere protection takes snapshots of your source virtual machine and replicates those snapshots to the target host where they are applied to the target virtual machine. The following diagram illustrates a typical agentless vSphere protection job. In this example, the replication interval is 15 minutes. Additionally, the example accounts for taking point-in-time snapshots of the target virtual machine. The example point-in-time snapshot interval is one hour. The job is started at 8:00.

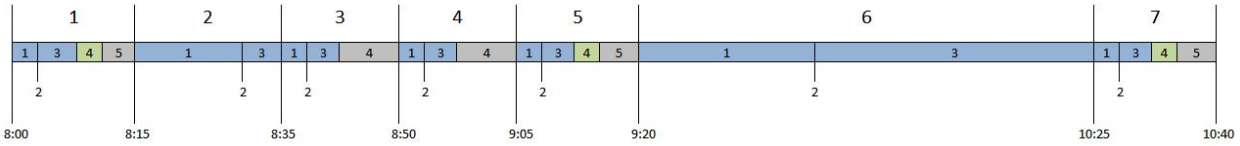

1. This is your first replication interval. It occurs between 8:00 and 8:15.
    1.1. Double-Take takes a replication snapshot of the source virtual machine.
    1.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run. Since this is your first replication interval and scheduled snapshots are enabled, a point-in-time snapshot of the target virtual machine will be triggered.
    1.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.
    1.4. Double-Take takes a point-in-time snapshot of the target virtual machine.
    1.5. Because the 15 minute replication interval has not been exhausted, Double-Take waits until the next replication interval is scheduled.
2. This is your second replication interval. It occurs between 8:15 and 8:35.
    2.1. Again, Double-Take takes a replication snapshot of the source virtual machine.
    2.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run. Because the point-in-time snapshot schedule of one hour has not elapsed, a point-in-time snapshot of the target virtual machine will not be triggered.
    2.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.

    In this replication interval, the process took longer than the 15 minutes allocated, so there is no delay before the next replication interval starts. Double-Take logs a message when this occurs.
3. This is your third replication interval. It occurs between 8:35 and 8:50.
    3.1. Double-Take takes a replication snapshot of the source virtual machine.
    3.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run, but it is not time yet.
    3.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.
    3.4. Double-Take waits until the next replication interval is scheduled.

4. This is your fourth replication interval. It occurs between 8:50 and 9:05.
    4.1. Double-Take takes a replication snapshot of the source virtual machine.
    4.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run, but it is not yet because it has only been 50 minutes since the job started.
    4.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.
    4.4. Double-Take waits until the next replication interval is scheduled.
5. This is your fifth replication interval. It occurs between 9:05 and 9:20.
    5.1. Double-Take takes a replication snapshot of the source virtual machine.
    5.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run. At this point, it has been one hour since the job started, so a point-in-time snapshot of the target virtual machine will be triggered.
    5.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.
    5.4. Double-Take takes a point-in-time snapshot of the target virtual machine.
    5.5. Double-Take waits until the next replication interval is scheduled.
6. This is your sixth replication interval. It occurs between 9:20 and 10:25.
    6.1. Double-Take takes a replication snapshot of the source virtual machine.
    6.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run. There is not because it has not been one hour since the last point-in-time snapshot.
    6.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.

    This replication interval took one hour and five minutes. Notice that the replication interval is set for every 15 minutes, but a new interval will not be started while a current interval is running.
7. This is your seventh replication interval. It occurs between 10:25 and 10:40.
    7.1. Double-Take takes a replication snapshot of the source virtual machine.
    7.2. Double-Take checks to see if a scheduled point-in-time snapshot is due to run. It has been one hour since the last point-in-time snapshot, so a new point-in-time snapshot of the target virtual machine will be triggered.
    7.3. Double-Take replicates and applies the replication snapshot of the source virtual machine to the target virtual machine.
    7.4. Double-Take takes a point-in-time snapshot of the target virtual machine.
    7.5. Double-Take waits until the next replication interval is scheduled.

Chapter 2 Agentless vSphere requirements

The configuration of an agentless vSphere job is different than other Double-Take jobs. Review the *Agentless vSphere overview* on page 4 first and then make sure your environment meets the following requirements.

- **ESXi hosts**—Your source and target hosts must be running one of the following ESXi versions. Note that this is ESXi that is supported. ESX is not supported.
  - ESXi 5.0 Update 2 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus
  - ESXi 5.1 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus
  - ESXi 5.5 Essentials, Essentials Plus, Standard, Enterprise, or Enterprise Plus
- **vCenter**—vCenter is not required, but if you are using it, then you must use version 5.0 or later.
- **Permissions**—If you want to limit the permissions required for the account that you will be using for your agentless vSphere job, your account must have at a minimum the permissions listed below. These permissions can be set at the vCenter, Datacenter, Cluster, or host level.
  - **Datastore**—Allocate space, Browse datastore, Low level file operations, Remove file, and Update virtual machine files
  - **Global**—Act as vCenter server and Licenses
  - **Host**, **Local Operations**—Create virtual machine, Delete virtual machine, Reconfigure virtual machine, and Relayout snapshots
  - **Network**—Assign Network
  - **Resource**—Assign virtual machine to resource pool
  - **Virtual Machine**—All virtual machine permissions are required.

  Make sure if you also define permissions at the VMs and Templates level in vCenter that you have not denied any of the required permissions listed above.
- **vMotion and DRS**
  - **Host vMotion**—Host vMotion is only supported if you are using vCenter.
  - **Storage vMotion**—Storage vMotion is only supported if you are using vCenter, with the exception of the caveats listed under DRS.
  - **DRS (Distributed Resource Scheduler)**—DRS is supported with Host vMotion and Storage vMotion, with the exception of the following caveats.
    - Due to a VMware issue with orphaned disk files, you should disable Automatic Storage vMotion using DRS.
    - Also due to the same issue, you should only manually vMotion after pausing your agentless vSphere jobs.

    See VMware Knowledge Base article 2055943 for more details on this issue.
- **ESXi and vCenter interoperability**—Any combination of ESXi and vCenter or ESXi standalone host is supported, with the following exceptions.
  - VMware does not support 5.1ESXi hosts on a 5.0 vCenter.
  - VMware does not support 5.5 ESXi hosts on a 5.0 vCenter.

    - Double-Take does not support the following source to target combations.
        - 5.1 ESXi standalone host to a 5.0 ESXi standalone host
        - 5.1 ESXi standalone host to a 5.0 ESXi host on a 5.0 vCenter
        - 5.1 ESXi host on a 5.1 vCenter to a 5.0 ESXi standalone host
        - 5.1 ESXi host on a 5.1 vCenter to a 5.0 ESXi host on a 5.0 vCenter
- **Source virtual machines**—Your source virtual machines must meet the following requirements and limitations.
    - **Hardware version**—Each virtual machine must be hardware version 8 or later. If your virtual machine is using hardware version 9, you will not be able to use ESXi version 5.0. This is a VMware limitation. If your virtual machine is using hardware version 10, you will not be able to use ESXi version 5.0 or 5.1. This is a VMware limitation.
    - **Operating system**—The guest operating system can be any operating system supported by VMware.
    - **Advanced network support, such as WAN support**—If you want Double-Take to automatically update networking on the replica during failover, for example for WAN support, you must have one of the following guest operating systems. If you choose not to have Double-Take automatically update networking on the replica during failover, you will have to update the network manually, but you can use any operating system supported by VMware.
        - **Windows**—Windows 2003, Windows 2003 R2, Windows 2008, Windows 2008 R2, Windows 2012, and Windows 2012 R2
        - **Linux**—Ubuntu 10.04 and 12.04, Redhat and CentOS 5 and 6, and SuSE 10 and 11
    - **Queiscing**—If you want to create quiescent replication snapshots (snapshots that are guaranteed to provide consistent application data), the following limitations apply.
        - The guest operating system must be Windows 2003 or later
        - VMware Tools must be installed on the guest
        - The guest must be powered on
    - **Raw device mapping**—You cannot use raw device mapping (RDM).
    - **Change block tracking**—Change block tracking must be supported on the source virtual machines, specifically the changeTrackingSupported flag must be set to true. See your VMware documentation for details on change block tracking.
    - **Independent disks**—Independent disks are not supported because VMware cannot take a snapshot of an independent disk, and Double-Take agentless vSphere replication technology is based on VMware snapshots.
    - **Controllers**—IDE controllers are not supported. You must use SCSI controllers.
- **Appliances**—Your controller appliance and your replication appliances will be an OVF (Open Virtualization Format) virtual machine included with Double-Take. You must install these appliances before you can establish protection. Keep in mind the following caveats for controller and replication appliances.
    - The appliances are pre-configured for optimal performance for most typical environments. While not required, you can modify memory, CPU, and other configurations to fit your specific environment.
    - You should not install or run anything else on these appliances. They must be dedicated to Double-Take processing only. Do not use them for any other activity (web server, database server, and so on).

- The firewall is disabled and should remain disabled.
- A single replication appliance can protect a maximum of 59 virtual disk files (.vmdk files). For example, if your virtual machines each have four disk files, you can protect 14 virtual machines (14*4=56). You may want to balance the replication among multiple replication appliances, depending on various factors such as the number of virtual machines you are protecting and the replication interval for your jobs.
- Because the controller appliance controls licensing for all of your replication appliances, the date and time (relative to UTC) on all of the appliances should be the same or within a few seconds of each other. Time differences could cause license invalidation.

- **Protocols and networking**—Your servers must meet the following protocol and networking requirements.
  - Your servers and appliances must have TCP/IP with static IP addressing.
  - IPv4 is the only supported version.
  - If you use DNS, you must have the appliances names and IP addresses in a forward lookup zone in DNS. Additionally, you must have the subnet the appliances will be using in a reverse lookup zone in DNS.
  - If you are not using DNS, you must configure the hosts file on each controller appliance and replication appliance to include mappings for all other replication appliances, controller appliances, vCenters, and ESXi servers.
- **Ports**—You must have TCP access for the following ports.
  - **22**—This port is for ssh (secure shell). You must have both inbound and outbound TCP access.
  - **443**—This port is required to add and manage the controller and replication appliances in vSphere. You must have outbound TCP access.
  - **5480**—This port is for accessing a controller or replication appliance using the web interface. You must have inbound TCP access.
  - **6320**—This port is for Java RMI (remote method invocation) communication for replication data between replication appliances. You must have both inbound and outbound TCP access.
  - **6325**—This port is for SOAP communication using http between the console and the controller appliance. You must have both inbound and outbound TCP access.
  - **6326**—This port is for SOAP communication using https between the controller appliance and the replication appliance. You must have both inbound and outbound TCP access.
- **Snapshots**—You can use third party tools to take snapshots of the source. Keep in mind the VMware snapshot limits for your version and that Double-Take requires one snapshot of the source, otherwise Double-Take replication will not work. Double-Take will delete its own snapshot on the source at the end of each replication interval.

  Third party snapshots of the replica on the target is not supported. You can take Double-Take point-in-time snapshots on the target if needed. Again, keep in mind the VMware snapshot limit for your version and that Double-Take requires at least two replication snapshots on the target in order for replication to work. Double-Take will account for the two replication snapshots by limiting the maximum number of point-in-time snapshots. Also Double-Take will delete any more than two replication snapshots on the target.
- **Supported configurations**—The following tables identify the supported host and virtual

machine configurations for an agentless vSphere job.

Host to Host Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect virtual machines on an active host to a standby host. The standby host has no active virtual machines, only the replica virtual machines. Replication will occur in only one direction.	X	
One to one active/active	You can protect virtual machines on an active host to another active host. Each host has active virtual machines in addition to replica virtual machines. Replication will be bi-directional between the hosts.	X	
Many to one	You can protect virtual machines on multiple hosts to one single host. This will consolidate your replica virtual machines to a single host.	X	
One to many	You can protect virtual machines on a single host to multiple target hosts, although individual virtual machines can only be protected once to a single target host. The will allow you to separate your replica virtual machines across multiple target hosts.	X	
Chained	You can protect virtual machines on host 1 to host 2, and you can protect virtual machines on host 2 to host 3, however, the virtual machines you are protecting on host 2 cannot be the replica virtual machines from host 1. They must be unique virtual machines that you are protecting.	X	
Single server	You can protect virtual machines within the same host or cluster.	X	
Standalone to standalone	Your hosts can be in a standalone to standalone configuration.	X	
Standalone to cluster	Your hosts can be in a standalone to cluster configuration.	X	
Cluster to standalone	Your hosts can be in a cluster to standalone configuration.	X	
Cluster to cluster	Your hosts can be in a cluster to cluster configuration.	X	

Virtual Machine to Virtual Machine Configuration	Description	Supported	Not Supported
One to one active/standby	You can protect a single virtual machine to a single, standby replica virtual machine. The replica virtual machine is not active while replication is occurring. The replica virtual machine will become active after failover.	X	
One to one active/active	You cannot protect a single virtual machine to a single, active replica virtual machine. The replica virtual machine is not active until after failover.		X
Many to one	You cannot protect a multiple virtual machines to a single virtual machine. Each virtual machine you protect creates only one replica virtual machine.		X
One to many	You cannot protect a single virtual machine to multiple virtual machines. A virtual machine can only be protected once, creating a single replica virtual machine.		X
Chained	You cannot protect a single virtual machine to a replica virtual machine to another replica virtual machine. The first replica virtual machine is not active, and therefore cannot be protected, until after failover.		X
Single server	You cannot protect a single virtual machine to itself.		X

Chapter 3 Creating an agentless vSphere job

Use these instructions to create an agentless vSphere job.

1. Make sure you have added your controller and replication appliances. See *Adding controller appliances* on page 57 and *Adding replication appliances* on page 58.
2. Click **Get Started** from the toolbar.
3. Select **Double-Take Availability** and click **Next**.
4. Select **Protect virtual servers on vSphere with agentless replication** and click **Next**.
5. Choose your source ESXi server. This is the ESXi server or cluster that contains the virtual machines that you want to protect. This is also where your source replication appliance is located. The list shows the ESXi servers or clusters currently associated with your controller appliance.

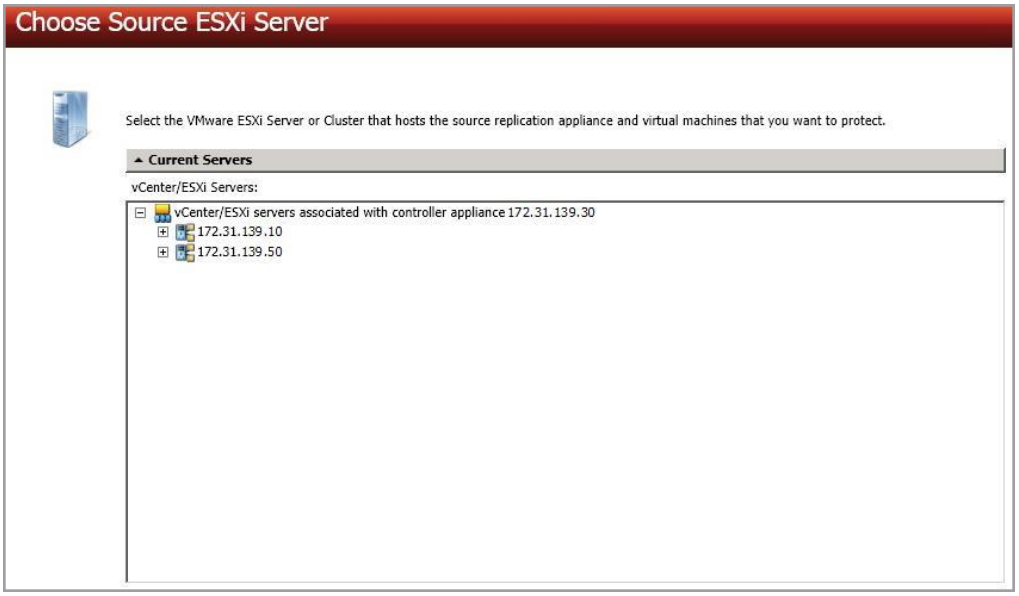

6. Click **Next** to continue.

7. Select the virtual machines on your source that you want to protect. A separate job will be created for each source that you select for protection. You can only select the number of virtual machines equal to the number of unused licenses you have available.

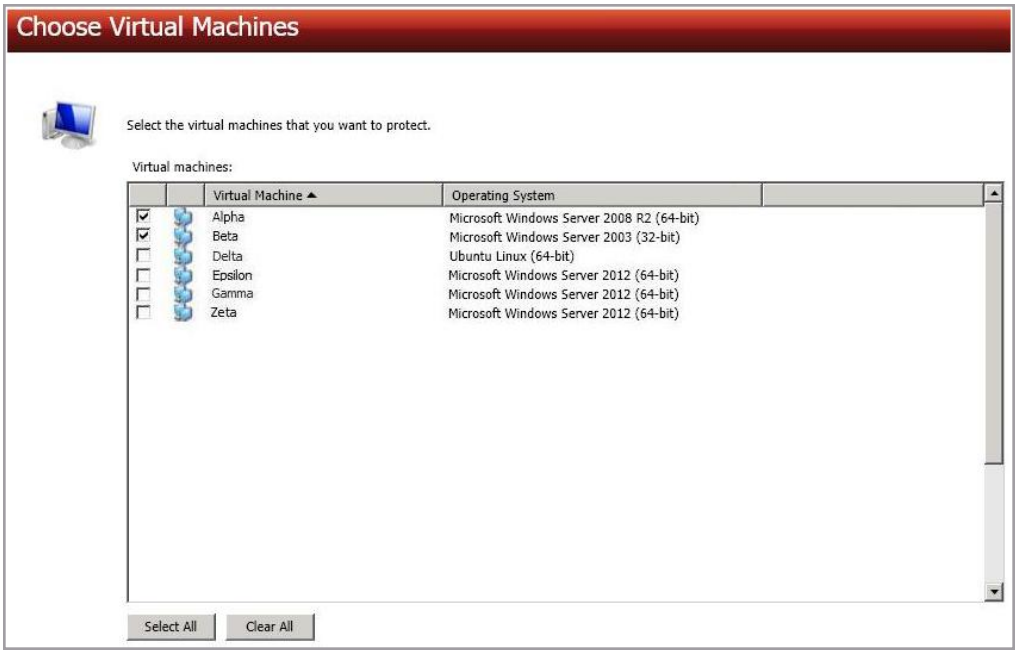

8. Click **Next** to continue.
9. Choose your target ESXi server. This is the ESXi server or cluster where the replica virtual machines will be hosted. This is also where your target replication appliance is located. The list shows the ESXi servers or clusters currently associated with your controller appliance.

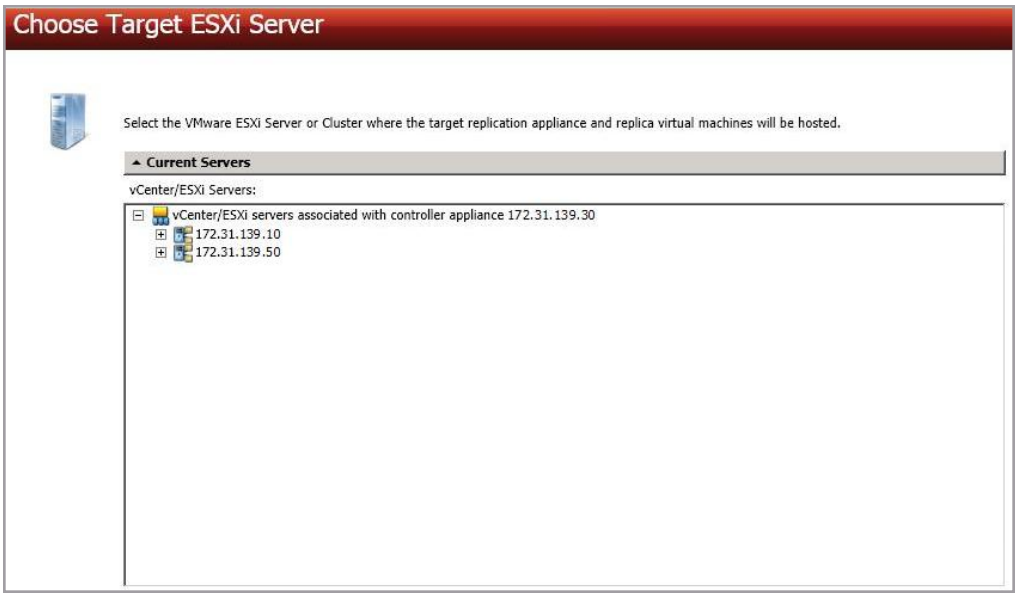

10. Click **Next** to continue.
11. You have many options available for your agentless vSphere job. Configure those options that are

applicable to your environment.

All agentless vSphere jobs will have the following sections available on the **Set Options** page.

- **Replica Virtual Machine Network Settings**
- **Virtual Switch Mappings**
- **Replication Settings**
- **Encryption**
- **Failover Options**
- **Snapshots**
- **Compression**

If you are protecting just one virtual machine, you will also have the following sections.

- **General**
- **Replica Virtual Machine Location**
- **Replica Host and Replication Appliances**

If you are protecting more than one virtual machine, you will have the **Virtual Machines Mappings** section instead, which is similar to the location and appliances sections.

As you can see, if you are protecting more than one virtual machine, there are a few settings that you will not have access to during job creation. In this case, default values will be used. You can modify the default values after the jobs have been created.

General

For the **Job name**, specify a unique name for your job.

Replica Virtual Machine Location

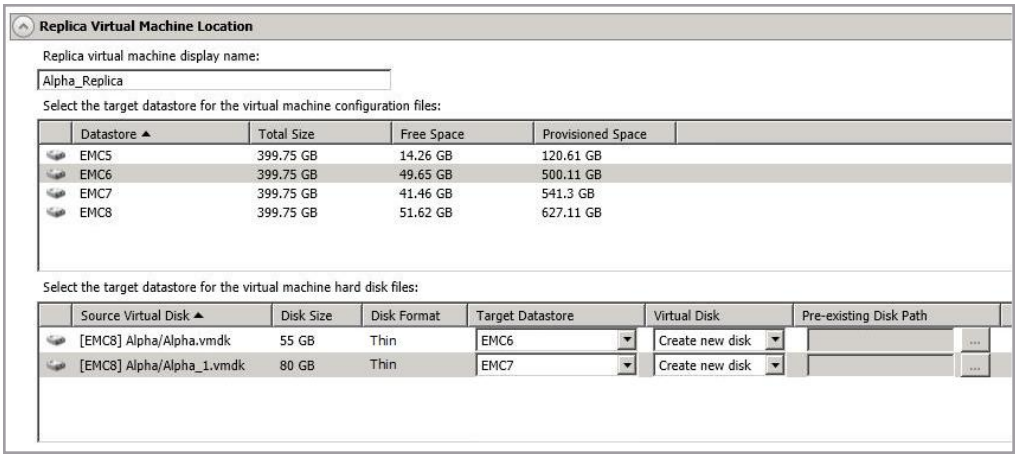

- **Replica virtual machine display name**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
- **Select the target datastore for the virtual machine configuration files**—Select one of the datastores from the list to indicate where on the target you want to store the replica virtual machine configuration files. Make sure the selected datastore has enough free space. The datastore you select here will become the selected datastore in the next table.
- **Target Datastore**—By default, the datastore you selected in the first table will be selected in the second table. If desired, specify a different datastore from the **Target Datastore** list to indicate where on the target you want to store each of the replica virtual machine disk files. Make sure the selected datastore has enough free space for the size of your source (plus growth) and the source's disk format (thick, thin, or flat). Additionally, if you are going to be enabling snapshots of the replica on the target, make sure you select a datastore that has enough space for the replica virtual machine disks and the snapshots.
- **Virtual Disk**—Specify if you want Double-Take to create a new disk for your replica virtual machine or if you want to use an existing disk.

  Reusing a virtual disk can be useful for pre-staging data on a LAN and then relocating the virtual disk to a remote site after the initial mirror is complete. You save time by skipping the virtual disk creation steps and performing a difference mirror instead of a full mirror. With pre-staging, less data will need to be sent across the wire initially. In order to use an existing virtual disk, it must be a valid virtual disk, it cannot be attached to any other virtual machine, and it cannot have any associated snapshots.

  Each pre-existing disk must be located on the target datastore specified. If you have copied the .vmdk file to this location manually, be sure you have also copied the associated -flat.vmdk file too. If you have used vCenter to copy the virtual machine, the associated file will automatically be copied. There are no restrictions on the file name of the .vmdk, but the associated -flat.vmdk file must have the same base name and the reference to that flat file in the .vmdk must be correct. Double-Take will move, not copy, the virtual disk files to the appropriate folders created by the replica, so make sure the selected target datastore is where you want the replica virtual disk to be located. Additionally, you cannot reuse a virtual disk if it has snapshots associated with it. You must delete all snapshots to reuse the virtual disk.

In a WAN environment, you may want to take advantage of using an existing disk by using a process similar to the following.

a. Create a job in a LAN environment, letting Double-Take create the virtual disk for you.
b. Complete the mirror process locally.
c. Delete the job and when prompted, do not delete the replica.
d. Delete all snapshots on the replica virtual machine.
e. Remove the replica virtual machine from the ESX inventory, which will delete the virtual machine configuration but will keep the associated .vmdk files.
f. Move the virtual disk files to the desired target datastore. Do not forget to move the associated -flat.vmdk file if you move the files manually.
g. Create a new protection job for the same source and reuse your existing disk.

If you have reused some existing disks and created some new disks, the numbering of the hard disks will not be identical on the source and the replica virtual machine. New disks will be created first and then existing disks will be attached. VMware assigns the hard disk numbers in order of creation and then those that are attached. The Virtual Device Node SCSI IDs will still be correct and there will be no impact within the guest of the replica virtual machine.

- **Pre-existing Disk Path**—This is the location on the selected **Target Datastore** of your existing virtual disks that you want to reuse.

Virtual Machines Mappings

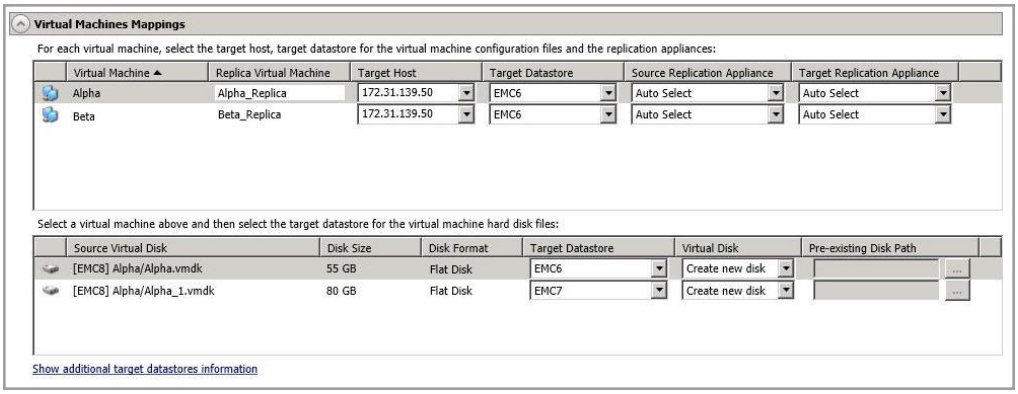

- **For each virtual machine, select the target host, target datastore for the virtual machine configuration files and the replication appliances**—The first table allows you to configure each virtual machine.
  - **Replica Virtual Machine**—Specify the name of the replica virtual machine. This will be the display name of the virtual machine on the host system.
  - **Target Host**—Select the specific server that you want to use, if your target is a cluster. If your target is not a cluster, this field will automatically be populated with your target.
  - **Target Datastore**—Specify the datastore on the target where you want to store the replica virtual machine configuration files. Make sure the selected datastore has enough free space. The datastore you select here will become the selected datastore in the next table.
  - **Source Replication Appliance**—You can specify a source replication appliance or select **Auto Select** to have Double-Take select a replication appliance for you based on the current load on all of your replication appliances.
  - **Target Replication Appliance**—You can specify a target replication appliance or select **Auto Select** to have Double-Take select a replication appliance for you based on the current load on all of your replication appliances.
- **Select a virtual machine above and then select the target datastore for the virtual machine hard disk files**—The second table is dynamically linked to the first. Select a virtual machine in the first table and configure the virtual disks for that virtual machine. Then select another virtual machine in the first table and configure the virtual disks for that virtual machine. Continue until you have configured the virtual disks for each virtual machine.
  - **Target Datastore**—By default, the datastore you selected in the first table will be selected in this table. If desired, specify a different datastore on the target where you want to store the replica virtual machine disk files. Make sure the selected datastore has enough free space for the size of your source (plus growth) and the source's disk format (thick, thin, or flat). Additionally, if you are going to be enabling snapshots of the replica on the target, make sure you select a datastore that has enough space for the replica virtual machine disks and the snapshots.
  - **Virtual Disk**—Specify if you want Double-Take to create a new disk for your replica virtual machine or if you want to use an existing disk.

Reusing a virtual disk can be useful for pre-staging data on a LAN and then relocating the virtual disk to a remote site after the initial mirror is complete. You save time by skipping the virtual disk creation steps and performing a difference mirror instead of a full mirror. With pre-staging, less data will need to be sent across the wire initially. In order to use an existing virtual disk, it must be a valid virtual disk, it cannot be attached to any other virtual machine, and it cannot have any associated snapshots.

Each pre-existing disk must be located on the target datastore specified. If you have copied the .vmdk file to this location manually, be sure you have also copied the associated -flat.vmdk file too. If you have used vCenter to copy the virtual machine, the associated file will automatically be copied. There are no restrictions on the file name of the .vmdk, but the associated -flat.vmdk file must have the same base name and the reference to that flat file in the .vmdk must be correct. Double-Take will move, not copy, the virtual disk files to the appropriate folders created by the replica, so make sure the selected target datastore is where you want the replica virtual disk to be located. Additionally, you cannot reuse a virtual disk if it has snapshots associated with it. You must delete all snapshots to reuse the virtual disk.

In a WAN environment, you may want to take advantage of using an existing disk by using a process similar to the following.

a. Create a job in a LAN environment, letting Double-Take create the virtual disk for you.
b. Complete the mirror process locally.
c. Delete the job and when prompted, do not delete the replica.
d. Delete all snapshots on the replica virtual machine.
e. Remove the replica virtual machine from the ESX inventory, which will delete the virtual machine configuration but will keep the associated .vmdk files.
f. Move the virtual disk files to the desired target datastore. Do not forget to move the associated -flat.vmdk file if you move the files manually.
g. Create a new protection job for the same source and reuse your existing disk.

If you have reused some existing disks and created some new disks, the numbering of the hard disks will not be identical on the source and the replica virtual machine. New disks will be created first and then existing disks will be attached. VMware assigns the hard disk numbers in order of creation and then those that are attached. The Virtual Device Node SCSI IDs will still be correct and there will be no impact within the guest of the replica virtual machine.

- **Pre-existing Disk Path**—This is the temporary location of your existing virtual disks that you want to reuse.

- **Show additional target datastores information**—Click this link to see storage information for the datastores on your target. This will help you select the appropriate datastore for your replica virtual machines.

Replica Virtual Machine Network Settings

The appearance and fields in this section will vary depending on if you have chosen to protect a single virtual machine or multiple virtual machines.

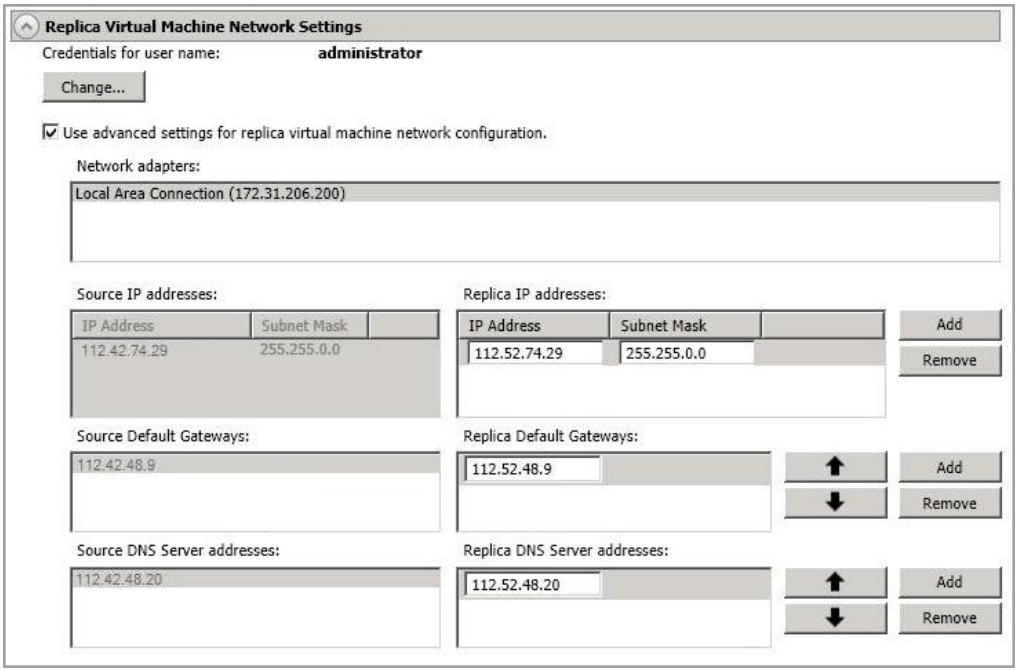

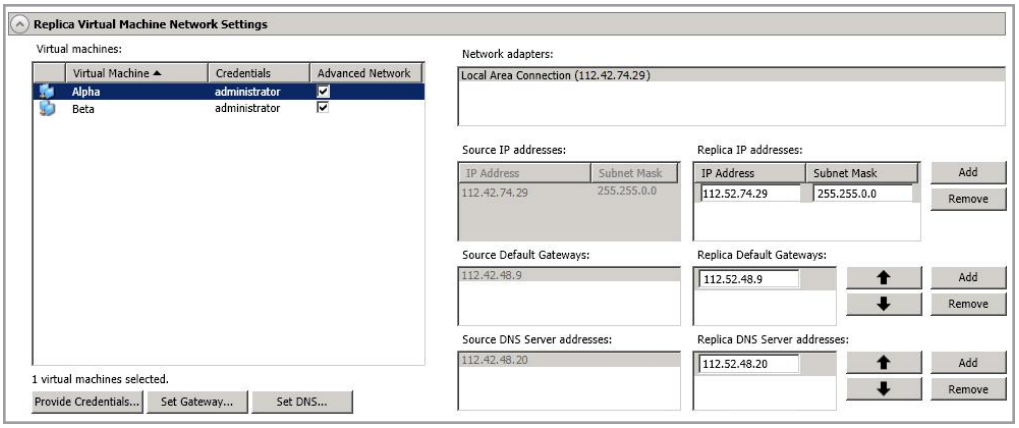

If your virtual machines are powered on and have VMware Tools version 8.3.8.9 (ESX 5.0 Update 1) or later installed, you can configure network settings for the replica virtual machines. This option is primarily used for WAN support.

- **Change**—If you have selected to protect a single virtual machine, click this button and specify credentials for the virtual machine.
- **Provide Credentials**—If you have selected to protect multiple virtual machines, you can select one or more virtual machines, click this button and specify credentials for the virtual machines. If necessary, use the credentials button again until all of the virtual machines that

you want to configure network settings for have credentials specified. You can right-click on any virtual machine in the list to access a **Select All** shortcut which will highlight all of the virtual machines in the list.

- **Use advanced settings for replica virtual machine network configuration**—Select this option to enable the replica virtual machine network setting configuration.
- **Advanced Network**—Select the virtual machines that you want to enable the replica virtual machine network setting configuration for.
- **Set Gateway**—If you are protecting multiple virtual machines, click this button to apply the same gateway settings that each source is using to its replica virtual machine. This is like a global update to quickly populate your replica virtual machine network settings. You can make additional modifications to each replica virtual machine's network settings if needed.
- **Set DNS**—If you are protecting multiple virtual machines, click this button to apply the same DNS settings that each source is using to its replica virtual machine. This is like a global update to quickly populate your replica virtual machine network settings. You can make additional modifications to each replica virtual machine's network settings if needed.
- **Network adapters**—Select a network adapter from the source and specify the **Replica IP addresses**, **Replica Default Gateways**, and **Replica DNS Server addresses** to be used after failover. If you add multiple gateways or DNS servers, you can sort them by using the arrow up and arrow down buttons. Repeat this step for each network adapter on the source.

Updates made during failover will be based on the network adapter name when protection is established. If you change that name, you will need to delete the job and re-create it so the new name will be used during failover.

If you update one of the advanced settings (IP address, gateway, or DNS server), then you must update all of them. Otherwise, the remaining items will be left blank. If you do not specify any of the advanced settings, the replica virtual machine will be assigned the same network configuration as the source.

By default, the source IP address will be included in the target IP address list as the default address. If you do not want the source IP address to be the default address on the target after failover, remove that address from the **Replica IP addresses** list.

Replica Host and Replication Appliances

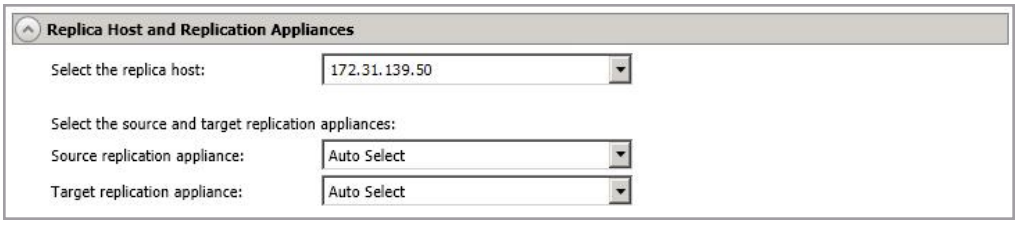

- **Select the replica host**—If you target is a cluster, select the specific server that you want to use. If your target is not a cluster, this field will automatically be populated with your target.
- **Source replication appliance**—You can specify a source replication appliance or select **Auto Select** to have Double-Take select a replication appliance for you based on the current load on all of your replication appliances.
- **Target replication appliance**—You can specify a target replication appliance or select **Auto Select** to have Double-Take select a replication appliance for you based on the current load on all of your replication appliances.

Virtual Switch Mappings

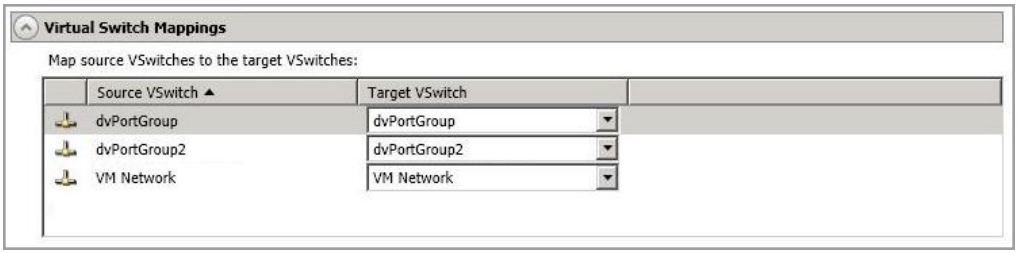

Map source VSwitches to the target VSwitches—Identify how you want to handle the network mapping after failover. Map each source host VSwitch to a target host VSwitch. If you reconfigure your virtual switches after you have created your job, you will need to edit the job to update your mappings, if you want to use the new configuration.

Replication Settings

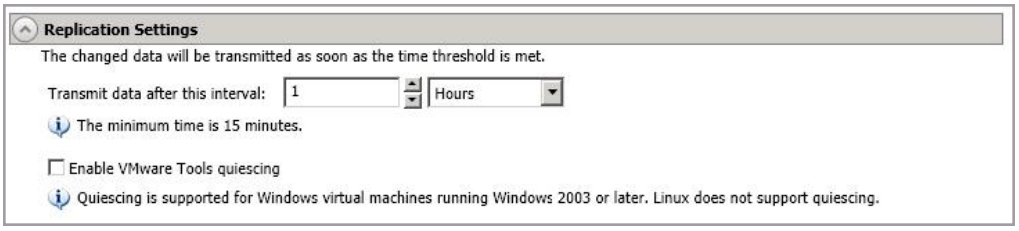

- **Transmit data after this interval**—Specify the interval (in days, hours, or minutes) for transmitting the replication snapshot. For more details on how the replication interval works, see *Replication process* on page 5.
- **Enable VMware Tools quiescing**—Select this option if you want quiescent replication snapshots (snapshots that are guaranteed to provide consistent application data). You must have VMware Tools installed on the source, your source must be powered on, and your source must be Windows 2003 or later. Linux operating systems do not support quiescing.

  Quiescing may impact performance on the source because disk operations will be suspended during the snapshot .VMware Tools instructs the applications running on the guest operating system that have registered for quiescing to suspend activity after they reach a good state. Depending on the type of applications you are running, for example a large transaction to a SQL server database, there may be a delay as applications quiesce, then the operating system quiesces, then the snapshot is taken. (Applications that are not registered for quiescing will have their disk write buffers flushed when the snapshot is taken.)

Encryption

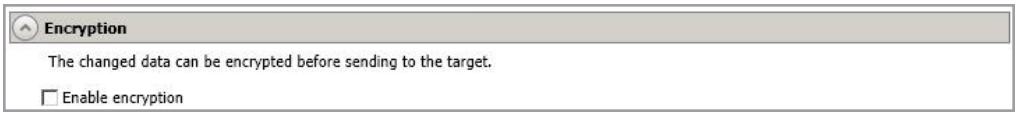

Enable encryption—Enable this option if you want to encrypt the mirroring and replication data before it is sent to the target.

Failover Options

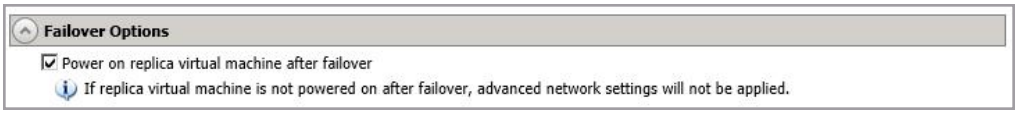

If you want the replica virtual machine to be powered on automatically after failover, enable **Power on replica virtual machine after failover**. The replica virtual machine must be on to apply any advanced network settings.

Snapshots

A snapshot is an image of the source replica data on the target taken at a single point in time. You can failover to a snapshot. However, you cannot access the snapshot to recover specific files or folders.

Turn on **Enable scheduled snapshots** if you want Double-Take to take snapshots automatically at set intervals.

- **Take snapshots on this interval**—Specify the interval (in days, hours, or minutes) for taking snapshots.
- **Maximum number of snapshots**—Because of VMware snapshot limitations and Double-Take replication processing, the maximum number of snapshots that you can retain is 29. At this point, the oldest snapshot will be deleted when a new snapshot needs to be created. If desired, you can decrease the number of snapshots that will be retained.

Snapshots are stored on the target datastore, so be sure that you have selected a datastore that has enough space for the replica virtual machine and the snapshots.

See *Replication process* on page 5 for details on how the snapshot of the source replica on the target works with the replication process.

Compression

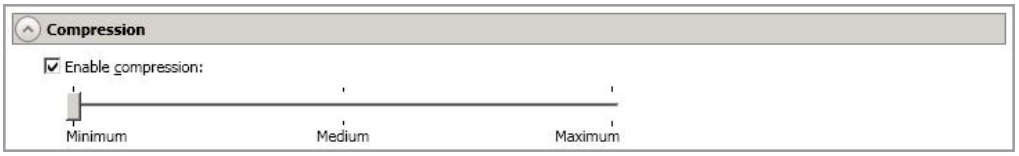

To help reduce the amount of bandwidth needed to transmit Double-Take data, compression allows you to compress data prior to transmitting it across the network. In a WAN environment this provides optimal use of your network resources. If compression is enabled, the data is compressed before it is transmitted from the source. When the target receives the compressed data, it decompresses it and then writes it to disk. You can set the level from **Minimum** to **Maximum** to suit your needs.

Keep in mind that the process of compressing data impacts processor usage on the source. If you notice an impact on performance while compression is enabled in your environment, either adjust to a lower level of compression, or leave compression disabled. Use the following guidelines to determine whether you should enable compression.

- If data is being queued on the source at any time, consider enabling compression.
- If the server CPU utilization is averaging over 85%, be cautious about enabling compression.
- The higher the level of compression, the higher the CPU utilization will be.
- Do not enable compression if most of the data is inherently compressed. Many image (.jpg, .gif) and media (.wmv, .mp3, .mpg) files, for example, are already compressed. Some images files, such as .bmp and .tif, are decompressed, so enabling compression would be beneficial for those types.
- Compression may improve performance even in high-bandwidth environments.
- Do not enable compression in conjunction with a WAN Accelerator. Use one or the other to compress Double-Take data.

All jobs from a single source connected to the same IP address on a target will share the same compression configuration.

12. Click **Next** to continue.
13. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

    Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

After a job is created, the results of the validation checks are logged to the job log. See the Double-Take *Reference Guide* for details on the various Double-Take log files.

14. Once your servers have passed validation and you are ready to establish protection, click **Finish**, and you will automatically be taken to the **Manage Jobs** page.

Chapter 4 Managing and controlling agentless vSphere jobs

The **Manage Jobs** page allows you to view status information about your jobs. You can also control your jobs from this page.

The jobs displayed in the right pane depend on the server group folder selected in the left pane. Every job for each server in your console session is displayed when the **Jobs on All Servers** group is selected. If you have created and populated server groups (see *Managing servers* on page 51), then only the jobs associated with the server or target servers in that server group will be displayed in the right pane.

Overview job information displayed in the top pane

The top pane displays high-level overview information about your jobs.

Column 1 (Blank)

The first blank column indicates the state of the job.

 The job is in a healthy state.

 The job is in a warning state. This icon is also displayed on any server groups that you have created that contain a job in a warning state.

 The job is in an error state. This icon is also displayed on any server groups that you have created that contain a job in an error state.

 The job is in an unknown state.

Job

The name of the job

Source Server

The name of the source. This could be the name or IP address of your source.

Target Server

The name of the target. This could be the name or IP address of your target.

Job Type

Each job type has a unique job type name. This job is an Agentless vSphere job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Activity

There are many different **Activity** messages that keep you informed of the job activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the job details. Keep in mind that **Idle** indicates console to server activity is idle, not that your servers are idle.

Mirror Status

- **In Progress**—Data is currently being mirrored.
- **Idle**—Data is not being mirrored.
- **Unknown**—The console cannot determine the status.

Replication Status

- **In Progress**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Unknown**—The console cannot determine the status.

Transmit Mode

- **Active**—Data is being transmitted to the target.
- **Unknown**—The console cannot determine the status.

Detailed job information displayed in the bottom pane

The details displayed in the bottom pane of the **Manage Jobs** page provide additional information for the job highlighted in the top pane. If you select multiple jobs, the details for the first selected job will be displayed.

Name

The name of the job

Target data state

This field is not applicable to agentless vSphere jobs.

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target. This field is only used when using pre-existing disks.

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This field is only used when using pre-existing disks.

Replication queue

This field is not applicable to agentless vSphere jobs.

Disk queue

This field is not applicable to agentless vSphere jobs.

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target. If you reset the target replication appliance, this value will be reset.

Bytes sent (compressed)

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as **Bytes sent**.

Connected since

This field is not applicable to agentless vSphere jobs.

Recent activity

Displays the most recent activity for the selected job, along with an icon indicating the success or failure of the last initiated activity. Click the link to see a list of recent activities for the selected job. You can highlight an activity in the list to display additional details about the activity.

Controller appliance

The controller appliance used by the replication appliances for this job

Source vCenter/ESXi host

The vCenter or ESXi server that the source is running on

Source virtual machine

The source virtual machine that is being protected

Target vCenter/ESXi host

The vCenter or ESXi server that the replica virtual machine is running on

Target datastore

The datastore on the target where the replica virtual machine is located

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Job controls

You can control your job through the toolbar buttons available on the **Manage jobs** page. If you select multiple jobs, some of the controls will apply only to the first selected job, while others will apply to all of the selected jobs. For example, **View Job Details** will only show details for the first selected job, while **Stop** will stop protection for all of the selected jobs.

If you want to control just one job, you can also right click on that job and access the controls from the pop-up menu.

Create a New Job 

This button leaves the **Manage Jobs** page and opens the **Get Started** page.

View Job Details 

This button leaves the **Manage Jobs** page and opens the **View Job Details** page.

Delete 

Stops (if running) and deletes the selected jobs.

If you no longer want to protect the source and no longer need the replica of the source on the target, select to delete the associated replica virtual machine. Selecting this option will remove the job and completely delete the replica virtual machine on the target.

If you no longer want to mirror and replicate data from the source to the target but still want to keep the replica of the source on the target, select to keep the associated replica virtual machine.

Provide Credentials 

Job credentials do not apply to agentless vSphere jobs.

View Recent Activity 

Displays the recent activity list for the selected job. Highlight an activity in the list to display additional details about the activity.

Start 

Starts the selected jobs.

Pause 

Agentless vSphere jobs cannot be paused.

Stop 

Stops the selected jobs. The jobs remain available in the console, but there will be no mirroring or replication data transmitted from the source to the target.

Take Snapshot 

Snapshots are not applicable to agentless vSphere jobs.

Manage Snapshots 

Snapshots are not applicable to agentless vSphere jobs.

Failover, Cutover, or Recover 

Starts the failover process. See *Failing over agentless vSphere jobs* on page 45 for the process and details of failing over an agentless vSphere job.

Failback 

Starts the failback process. Failback does not apply to agentless vSphere jobs.

Restore 

Starts the restoration process. Restore does not apply to agentless vSphere jobs.

Reverse 

Reverses protection. Reverse protection does not apply to agentless vSphere jobs.

Recover 

Recovers the selected DR job. Recovery does not apply to agentless vSphere jobs.

Undo Failover or Cutover 

Cancels failover by undoing it. Undo does not apply to agentless vSphere jobs.

View Job Log 

Opens the job log. On the right-click menu, this option is called **View Logs**, and you have the option of opening the job log, source server log (which is the source replication appliance), or target server log (which is the target replication appliance).

If you need to manually collect agentless vSphere job logs, perform the following steps on each controller and replication appliance.

1. Connect to the appliance using an ssh client, like Putty.
2. Change to the doubletake directory using cd /opt/visionsolutions/doubletake.

3. Run the following command (exactly as below without substitutions) to create an archive called server.tgz in the doubletake directory.

   tar czvf $(hostname).tgz bin data etc log

4. You can then copy (using a tool like WinSCP) or FTP the file to another location, if needed.

Other Job Actions 

Opens a small menu of other job actions. These job actions are not available for agentless vSphere jobs.

Filter

Select a filter option from the drop-down list to only display certain jobs. You can display **Healthy jobs**, **Jobs with warnings**, or **Jobs with errors**. To clear the filter, select **All jobs**. If you have created and populated server groups, then the filter will only apply to the jobs associated with the server or target servers in that server group. See *Managing servers* on page 51.

Type a server name

Displays only jobs that contain the text you entered. If you have created and populated server groups, then only jobs that contain the text you entered associated with the server or target servers in that server group will be displayed. See *Managing servers* on page 51.

Overflow Chevron 

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Viewing agentless vSphere job details

From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.

Review the following table to understand the detailed information about your job displayed on the **View Job Details** page.

Job name

The name of the job

Job type

Each job type has a unique job type name. This job is an Agentless vSphere job. For a complete list of all job type names, press F1 to view the Double-Take Console online help.

Health

 The job is in a healthy state.

 The job is in a warning state.

 The job is in an error state.

 The job is in an unknown state.

Activity

There are many different **Activity** messages that keep you informed of the job activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the rest of the job details.

Connection ID

The incremental counter used to number connections. The number is incremented when a connection is created. It is also incremented by internal actions, such as an auto-disconnect and auto-reconnect. The lowest available number (as connections are created, stopped, deleted, and so on) will always be used. The counter is reset to one each time the Double-Take service is restarted.

Transmit mode

- **Active**—Data is being transmitted to the target.
- **Unknown**—The console cannot determine the status.

Target data state

This field is not applicable to agentless vSphere jobs.

Target route

This field is not applicable to agentless vSphere jobs.

Compression

- **On** / **Level**—Data is compressed at the level specified.
- **Off**—Data is not compressed.

Encryption

- **On**—Data is being encrypted before it is sent from the source to the target.
- **Off**—Data is not being encrypted before it is sent from the source to the target.

Bandwidth limit

This field is not applicable to agentless vSphere jobs.

Connected since

This field is not applicable to agentless vSphere jobs.

Additional information

Depending on the current state of your job, you may see additional information displayed to keep you informed about the progress and status of your job. If there is no additional information, you will see (None) displayed.

Mirror status

- **In Progress**—Data is currently being mirrored.
- **Idle**—Data is not being mirrored.
- **Unknown**—The console cannot determine the status.

Mirror percent complete

The percentage of the mirror that has been completed

Mirror remaining

The total number of mirror bytes that are remaining to be sent from the source to the target

Mirror skipped

The total number of bytes that have been skipped when performing a difference. These bytes are skipped because the data is not different on the source and target. This field is only used when using pre-existing disks.

Replication status

- **In Progress**—Data is being replicated to the target.
- **Ready**—There is no data to replicate.
- **Unknown**—The console cannot determine the status.

Replication queue

This field is not applicable to agentless vSphere jobs.

Disk queue

This field is not applicable to agentless vSphere jobs.

Bytes sent

The total number of mirror and replication bytes that have been transmitted to the target

Bytes sent compressed

The total number of compressed mirror and replication bytes that have been transmitted to the target. If compression is disabled, this statistic will be the same as **Bytes sent**.

Validating an agentless vSphere job

Over time, you may want to confirm that any changes in your network or environment have not impacted your Double-Take job. Use these instructions to validate an existing job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Validate job properties**.
3. Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.
4. Once your servers have passed validation, click **Close**.

Editing an agentless vSphere job

Use these instructions to edit an agentless vSphere job.

1. From the **Manage Jobs** page, highlight the job and click **View Job Details** in the toolbar.
2. In the **Tasks** area on the right on the **View Job Details** page, click **Edit job properties**.
3. You will see the same options for your agentless vSphere job as when you created the job, but you will not be able to edit all of them. If desired, edit those options that are configurable for an existing job. See *Creating an agentless vSphere job* on page 12 for details on each job option.

   

   Changing some options may require Double-Take to automatically disconnect, reconnect, and remirror the job.

4. Click **Next** to continue.
5. Double-Take validates that your source and target are compatible. The **Summary** page displays your options and validation items.

   Errors are designated by a white X inside a red circle. Warnings are designated by a black exclamation point (!) inside a yellow triangle. A successful validation is designated by a white checkmark inside a green circle. You can sort the list by the icon to see errors, warnings, or successful validations together. Click on any of the validation items to see details. You must correct any errors before you can continue. Depending on the error, you may be able to click **Fix** or **Fix All** and let Double-Take correct the problem for you. For those errors that Double-Take cannot correct automatically, you will need to modify the source or target to correct the error, or you can select a different target. You must revalidate the selected servers, by clicking **Recheck**, until the validation check passes without errors.

   After a job is created, the results of the validation checks are logged to the job log. See the Double-Take *Reference Guide* for details on the various Double-Take log files.
6. Once your servers have passed validation and you are ready to update your job, click **Finish**.

Viewing an agentless vSphere job log

You can view a job log file through the Double-Take Console by selecting **View Job Log** from the toolbar on the **Manage Jobs** page. Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple jobs. When the Double-Take Console is closed, all logging windows will automatically close.

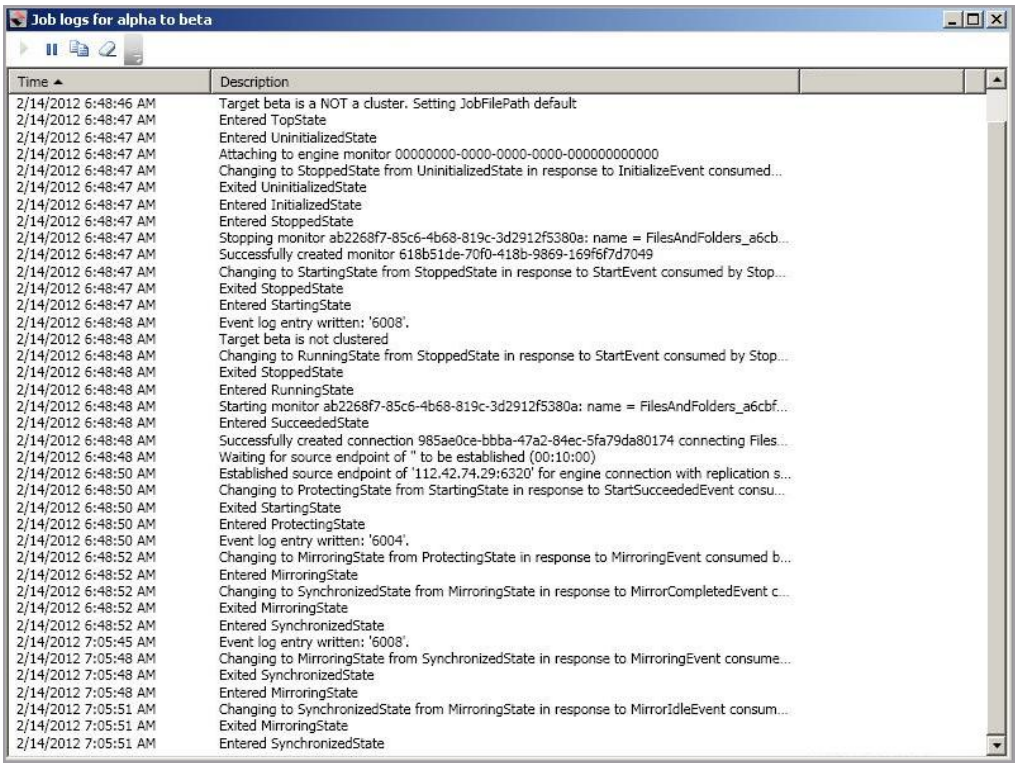

The following table identifies the controls and the table columns in the **Job logs** window.

Start

This button starts the addition and scrolling of new messages in the window.

Pause

This button pauses the addition and scrolling of new messages in the window. This is only for the **Job logs** window. The messages are still logged to their respective files on the server.

Copy

This button copies the messages selected in the **Job logs** window to the Windows clipboard.

Clear

This button clears the **Job logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Job logs** window.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Updating to a new controller appliance

Use the following process to update an existing job to a new controller appliance.

1. Deploy a new controller appliance using a new name and IP address.
2. Make sure your old controller appliance is powered off.
3. Remove the old controller appliance from the Double-Take Console using the **Remove Server** button on the **Manage Servers** page.
4. Add the new controller appliance to the Double-Take Console. See *Adding controller appliances* on page 57.
5. Complete a host transfer of the license from the old controller appliance, adding that transferred license to the new controller appliance. See the *Double-Take Installation, Licensing, and Activation* document for details on a host transfer.
6. Associate your replication appliances with the new controller appliance. See *Adding replication appliances* on page 58.

After you have associated your replication appliances with your new controller appliance, your existing job will appear in the console.

Failing over agentless vSphere jobs

When you are ready to failover, you can start it manually from the console.

1. On the **Manage Jobs** page, highlight the job that you want to failover and click **Failover, Cutover, or Recover** in the toolbar.
2. Select the type of failover to perform.
   - **Failover to live data**—Select this option to initiate a full, live failover using the last complete replication cycle. This option will shutdown the source machine (if it is online), stop the protection job, and start the replica virtual machine on the target with full network connectivity.
   - **Perform test failover**—This option is not available for agentless vSphere jobs.
   - **Failover to a snapshot**—Select this option to initiate a full, live failover without using the last replication cycle. Instead, select a snapshot and the data on the target will be reverted to that snapshot. This option will not be available if there are no snapshots on the target.
3. Normally, you would select how you want to handle the data in the target queue. However, these options are not currently available.
4. When you are ready to begin failover, click **Failover**.

Chapter 5 Double-Take Console

After you have installed the console, you can launch it by selecting **Double-Take**, **Double-Take Console** from your **Programs**, **All Programs**, or **Apps**, depending on your operating system.

The Double-Take Console is used to protect and monitor your servers and jobs. Each time you open the Double-Take Console, you start at the **Home** page. This page provides a high-level overview of the status of your jobs.

The appearance of the **Home** page is the same for all users. However, other console pages may have variances in the appearance depending on the Double-Take products that you have installed, the Double-Take license keys on your servers, and the type of job you are working with.

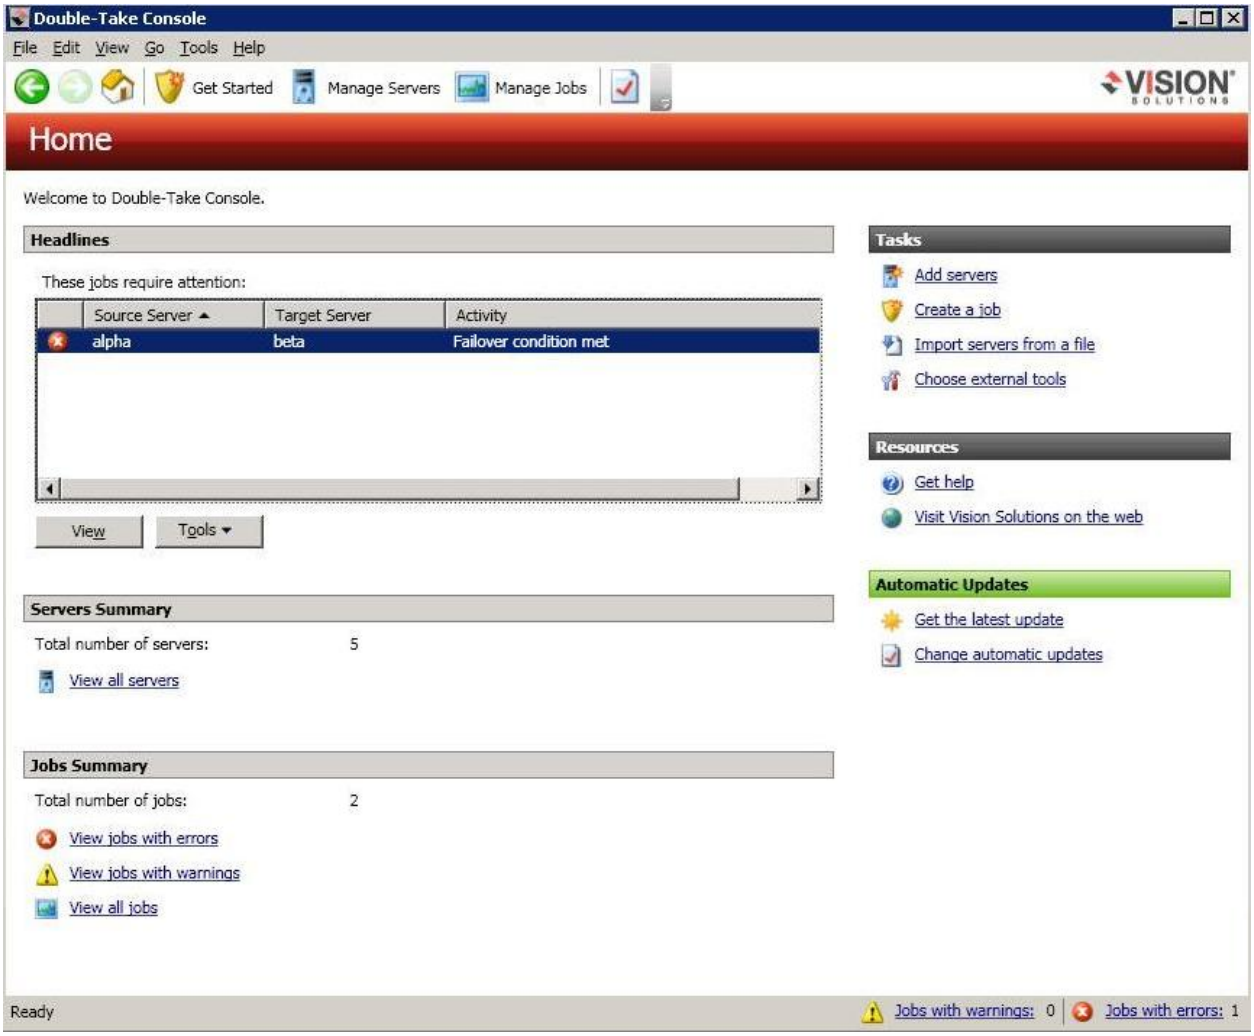

- **Headlines**—The top section gives a quick overview of any jobs that require attention as well as providing quick access buttons.
  - **These jobs require attention**—Any jobs that require attention (those in an error state) are listed. You will see the source and target server names listed, as well as a short description of the issue that requires your attention. If the list is blank, there are no jobs that

require immediate attention.

- **View**—If you highlight a job in the list and click **View**, you will go to the **View Job Details** page where you can see more detailed information about the job.
- **Tools**—Select this drop-down list to launch other Vision Solutions consoles.

- **Servers Summary**—The middle section summarizes the servers in your console.
  - **Total number of servers**—This field displays the number of servers that you have been added to the console.
  - **View all servers**—Select this link to go to the **Manage Servers** page where you can view, edit, add, remove, or manage the servers in your console. See *Managing servers* on page 51.
- **Jobs Summary**—The bottom section summarizes the jobs in your console.
  - **Total number of jobs**—This field displays the number of jobs running on the servers in your console.
  - **View jobs with errors**—Select this link to go to the **Manage Jobs** page, where the **Filter: Jobs with errors** will automatically be applied.
  - **View jobs with warnings**—Select this link to go to the **Manage Jobs** page, where the **Filter: Jobs with warnings** will automatically be applied.
  - **View all jobs**—Select this link to go to the **Manage Jobs** page and view all jobs.

At the bottom of the Double-Take Console, you will see a status bar. At the right side, you will find links for **Jobs with warnings** and **Jobs with errors**. This lets you see quickly, no matter which page of the console you are on, if you have any jobs that need your attention. Select this link to go to the **Manage Jobs** page, where the appropriate **Filter: Jobs with warnings** or **Filter: Jobs with errors** will automatically be applied.

Double-Take Console requirements

You must meet the following requirements for the Double-Take Console.

- **Operating system**—The Double-Take Console can be run from a Windows source or target. It can also be run from a 32-bit or 64-bit physical or virtual machine running Windows 8, Windows 7, Windows Vista, or Windows XP Service Pack 2 or later.
- **Microsoft .NET Framework**—Microsoft .NET Framework version 4.0 Update 3 or later is required. (The full .NET 4.0.3 is required, not just the Client Profile.)
- **Screen resolution**—For best results, use a 1024x768 or higher screen resolution.

The Double-Take installation prohibits the console from being installed on Server Core. Because Windows 2012 allows you to switch back and forth between Server Core and a full installation, you may have the console files available on Server Core, if you installed Double-Take while running in full operating system mode. In any case, you cannot run the Double-Take Console on Server Core.

Console options

There are several options that you can set that are specific to the Double-Take Console. To access these console options, select **Options** from the toolbar.

- **Monitoring**—This section is used to determine how the console monitors your Double-Take servers.
    - **Monitoring interval**—Specifies how often, in seconds, the console refreshes the monitoring data. The servers will be polled at the specified interval for information to refresh the console.
    - **Automatic retry**—This option will have the console automatically retry server login credentials, after the specified retry interval, if the server login credentials are not accepted. Keep in mind the following caveats when using this option.
        - This is only for server credentials, not job credentials.
        - A set of credentials provided for or used by multiple servers will not be retried for the specified retry interval on any server if it fails on any of the servers using it.
        - Verify your environment's security policy when using this option. Check your policies for failed login lock outs and resets. For example, if your policy is to reset the failed login attempt count after 30 minutes, set this auto-retry option to the same or a slightly larger value as the 30 minute security policy to decrease the chance of a lockout.
        - Restarting the Double-Take Console will automatically initiate an immediate login.
        - Entering new credentials will initiate an immediate login using the new credentials.
    - **Retry on this interval**—If you have enabled the automatic retry, specify the length of time, in minutes, to retry the login.
- **Server Communication**—This section is used to determine how the console communicates with your Double-Take servers.
    - **Default port for XML web services protocol**—Specifies the port that the console will use when sending and receiving data to Double-Take servers. By default, the port is 6325. Changes to the console port will not take effect until the console is restarted.
    - **Default port for legacy protocol**—If you are using an older Double-Take version, you will need to use the legacy protocol port. This applies to Double-Take versions 5.1 or earlier.
- **Diagnostics**—This section assists with console troubleshooting.
    - **Export Diagnostic Data**—This button creates a raw data file that can be used for debugging errors in the Double-Take Console. Use this button as directed by technical support.
    - **View Log File**—This button opens the Double-Take Console log file. Use this button as directed by technical support. You can also select **View**, **View Console Log File** to open the Double-Take Console log file.
    - **View Data File**—This button opens the Double-Take Console data file. Use this button as directed by technical support. You can also select **View**, **View Console Data File** to open the Double-Take Console data file.
- **Automatic Updates**—This section is for automatically updating your console.

- **Automatically check for updates**—By default, each time the console is started, it will automatically check the Vision Solutions web site to see if there is updated console software available. If there is updated console software available, an **Automatic Updates** section will appear on the **Home** page. Click **Get the latest update** to download and install the updated console software.

  If you want to disable the automatic check for updates, click **Change automatic updates** or select **Options** from the toolbar. On the **Options** page, deselect **Automatically check for updates** to disable the automatic check.

  You can also manually check for updates by selecting **Help**, **Check for Updates**.

  - **Update available**—If there is an update available, click **Get Update**. The dialog box will close and your web browser will open to the Vision Solutions web site where you can download and install the update.
  - **No update available**—If you are using the most recent console software, that will be indicated. Click **Close**.
  - **No connection available**—If the console cannot contact the update server of if there is an error, the console will report that information. The console log contains a more detailed explanation of the error. Click **Check using Browser** if you want to open your browser to check for console software updates. You will need to use your browser if your Internet access is through a proxy server.

- **License Inventory**—This section controls if the console contains a license inventory. This feature may not appear in your console if your service provider has restricted access to it.

  - **Enable license inventory**—This option allows you to use this console to manage the Double-Take licenses assigned to your organization. When this option is enabled, the **Manage License Inventory** page is also enabled.

- **Default Installation Options**—These fields are not used for agentless vSphere protection.
- **Default Windows Installation Options**—These fields are not used for agentless vSphere protection.
- **Default Linux Installation Options**—These fields are not used for agentless vSphere protection.

Chapter 6 Managing servers

To manage the servers in your console, select **Manage Servers** from the toolbar. The **Manage Servers** page allows you to view, edit, add, remove, or manage the servers in your console.

You can also organize the servers that are in your console into groups, allowing you to filter the servers you are viewing based on your organization. The servers displayed in the right pane depend on the server group folder selected in the left pane. Every server in your console session is displayed when the **All Servers** group is selected. If you have created and populated server groups under **My Servers**, then only the servers in the selected group will displayed in the right pane.

If you have uninstalled and reinstalled Double-Take on a server, you may see the server twice on the **Manage Servers** page because the reinstall assigns a new unique identifier to the server. One of the servers (the original version) will show with the red X icon. You can safely remove that server from the console.

Right pane display

The following table identifies the columns displayed in the right pane of the **Manage Servers** page.

Column 1 (Blank)

The first blank column indicates the machine type.

 Double-Take source or target server which could be a physical server, virtual machine, or a cluster node

 Double-Take source or target server which is a Windows cluster

 VMware server which could be a vCenter server or an ESX or ESXi host.

 Double-Take controller appliance

 Double-Take replication appliance

 Double-Take Reporting Service server

 Offline server which means the console cannot communicate with this machine.

 Server error which means the console can communicate with the machine, but it cannot communicate with Double-Take on it.

Column 2 (Blank)

The second blank column indicates the security level

Processing—The console is attempting to communicate with machine.

Administrator access—This level grants full control.

Monitor only access—This level grants monitoring privileges only.

No security access—This level does not allow monitoring or control.

Server

The name or IP address of the server. If you have specified a reserved IP address, it will be displayed in parenthesis.

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the server details. See *Viewing server details* on page 61.

Version

The product version information

Licensing Status

The status of the license on the server. If your license is expired, any jobs using that server will be in an error state.

Product

The Double-Take products licensed for the server or the Double-Take role for the server.

License Key

The license keys associated with the products licensed for the server. If your license is not valid for the operating system on your server, the license key will be identified as Invalid License Key. There will be no license key listed for those servers that are not licensed, like a VMware server.

Serial Number

The serial number associated with the license key

Main toolbar and right-click menu

The following options are available on the main toolbar of the **Manage Servers** page and the right-click menu. Some of the options are only available in the right-click menu. Some options are only available for a single selected server and others are available for multiple selected servers.

Add Servers 

Adds a new server. This button leaves the **Manage Servers** page and opens the **Add Servers** page. See *Adding controller appliances* on page 57.

Add Replication Appliance 

Adds a new replication appliance. You must have a controller appliance already inserted in your console before you can add any replication appliances. This button leaves the **Manage Servers** page and opens the **Add Replication Appliances** page. You can also add replication appliances by editing a controller appliance's server properties. See *Controller appliance associations properties* on page 64.

View Server Details 

Views detailed information about a server. This button leaves the **Manage Servers** page and opens the **View Server Details** page. See *Viewing server details* on page 61.

Remove Server 

Removes the server from the console. For agentless vSphere jobs, you cannot remove replication appliances from the **Manage Servers** page. You must remove the replication appliance from the controller appliance's server properties (see *Controller appliance associations properties* on page 64) or from the Add Replication Appliance page (see *Adding replication appliances* on page 58). Keep in mind, if you remove a controller appliance from the console without removing the replication appliance associations, the replication appliances will still be associated with that controller appliance if you add the controller appliance back into this or into another console.

Provide Credentials 

Changes the login credentials that the Double-Take Console use to authenticate to a server. This button opens the **Provide Credentials** dialog box where you can specify the new account information. See *Providing server credentials* on page 60. You will remain on the **Manage Servers** page after updating the server credentials.

Manage Group Assignments 

Allows you to assign, move, and remove the selected server from specific server groups. This buttons opens the Manage Group Assignments dialog box where you can assign and unassign the server to specific server groups. The server will appear in server groups marked with a checkmark, and will not appear in groups without a checkmark. Servers assigned to a server group will automatically appear in parent server groups.

Install 

Installs or upgrades Double-Take on the selected server. This option is not applicable to agentless vSphere protection.

Uninstall 

Uninstalls Double-Take on the selected server. This option is not applicable to agentless vSphere protection.

Copy 

Copies the information for the selected servers. You can then paste the server information as needed. Each server is pasted on a new line, with the server information being comma-separated.

Paste 

Pastes a new-line separated list of servers into the console. Your copied list of servers must be entered on individual lines with only server names or IP addresses on each line.

View Server Events 

Views event messages for a server. The option is not applicable to agentless vSphere protection.

View Server Logs 

Views the Double-Take logs messages for a server. This button opens the **Logs** window. This separate window allows you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

Activate Online 

Activates licenses and applies the activation keys to servers in one step. You must have Internet access for this process. You will not be able to activate a license that has

already been activated but has not been deactivated. In that case, you will be prompted to complete a host transfer. Ideally, you should deactivate the license instead of doing a host transfer See the Double-Take Console online help for details on deactivating licenses and host transfers.

Gather Support Diagnostics 

Executes the diagnostic DTInfo utility which collects configuration data for use when reporting problems to technical support. It gathers Double-Take log files; Double-Take and system settings; network configuration information such as IP, WINS, and DNS addresses; and other data which may be necessary for technical support to troubleshoot issues. You will be prompted for a location to save the resulting file which is created with the information gathered. Because this utility is gathering several pieces of information, across the network to your console machine, it may take several minutes to complete the information gathering and sending the resulting file to the console machine.

This option is disabled for replication appliances. That data will be collected when data is collected for a controller appliance.

View Replication Service Details 

Views the replication service details for a server. This option is not applicable to agentless vSphere servers.

Refresh 

Refreshes the status of the selected servers.

Overflow Chevron 

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Left pane toolbar

Between the main toolbar and the left pane is a smaller toolbar. These toolbar options control the server groups in the left pane.

Create New Server Group 

Creates a new server group below the selected group

Rename Server Group 

Allows you to rename the selected server group

Delete Server Group 

Deletes the selected server group. This will not delete the servers in the group, only the group itself.

Overflow Chevron 

Displays any toolbar buttons that are hidden from view when the window size is reduced.

Adding controller appliances

The first time you start the console, the **Manage Servers** page is empty. In order to protect and monitor your virtual machines, you must insert your controller appliance in the console.

1. Select **Get Started** from the toolbar.
2. Select **Add servers** and click **Next**.
3. On the **Manual Entry** tab, specify the controller appliance information.
   - **Server**—This is the name or IP address of the controller appliance to be added to the console.
   - **User name**—Specify the root user or another user that has the administrator role on the appliance.
   - **Password**—Specify the password associated with the **User name** you entered.
   - **Domain**—If you are working in a domain environment, specify the **Domain**.
4. After you have specified the appliance information, click **Add**.
5. Do not insert your replication appliances. The replication appliances will be inserted into the console using the Add Replication Appliances pages. See *Adding replication appliances* on page 58.
6. If you need to remove appliances from the list of **Servers to be added**, highlight it and click **Remove**. You can also remove all of them with the **Remove All** button.
7. When your list of **Servers to be added** is complete, click **OK**.

Adding replication appliances

You can add a replication appliance by selecting **Get Started**, **Add replication appliances** or on the **Manage Servers** page, click **Add Replication Appliance** from the toolbar.

From the **Add Replication Appliances** page, you can associate replication appliances to controller appliances.

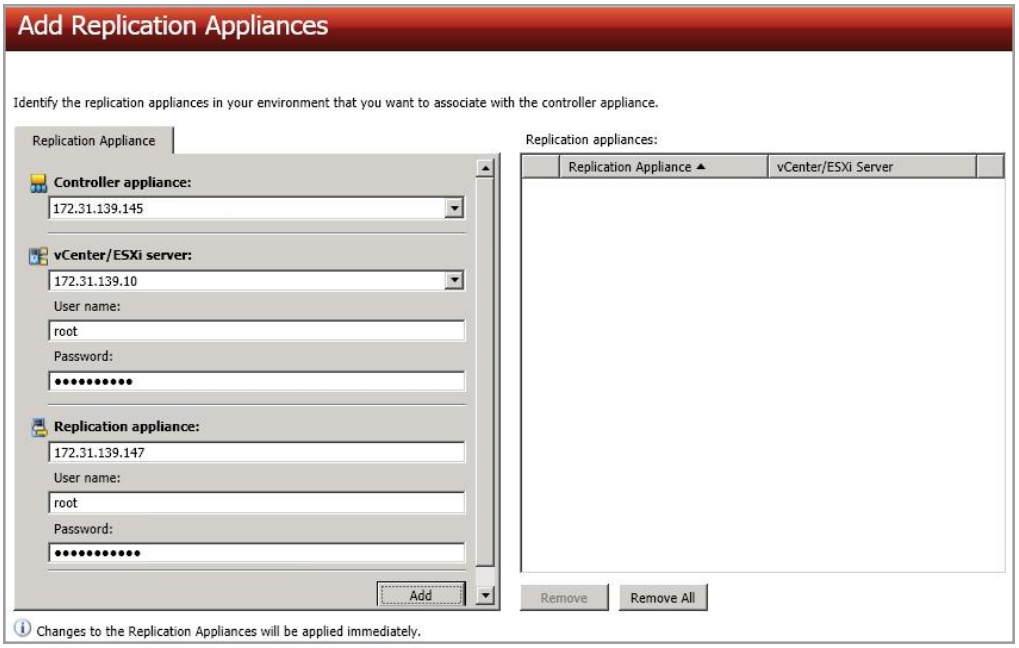

- **Controller appliance**—Select a controller appliance from the list. The replication appliances you add on this page will be associated with the controller appliance you select.
- **vCenter/ESXi server**—Specify the IP address of the vCenter or ESXi server that is hosting the replication appliance the selected controller appliance will communicate with. If you are using vCenter, specify your vCenter. Only specify your ESXi server if it is a standalone server. If you are not using the default vCenter/ESXi port, you can specify the port number that you are using by entering IP_address:port_number. For example, you might enter 172.31.139.10:1124.
    - **User name**—Specify the root user or another user that has the administrator role on the specified server.
    - **Password**—Specify the password associated with the **User name** that you entered.
- **Replication appliance**—Specify the IP address of the replication appliance.
    - **User name**—Specify the root user or another user that has the administrator role on the specified replication appliance.
    - **Password**—Specify the password associated with the **User name** that you entered.

After you have configured your replication appliance to controller appliance association, click **Add**. Repeat the process to associate additional replication appliances to your controller appliance. If you need to remove any replication appliance associations from a controller appliance, highlight the appliance in the list and click **Remove**. You can also remove all of the associated replication appliances by clicking **Remove All**. Keep in mind that you cannot remove appliance associations if a job is using the

appliance. You can also manage your appliance associations through the controller appliance's server properties. See *Controller appliance associations properties* on page 64.

Providing server credentials

To update the security credentials used for a specific server, select **Provide Credentials** from the toolbar on the **Manage Servers** page. When prompted, specify the **User name**, **Password**, and **Domain** of the account you want to use for this server. Click **OK** to save the changes.

This process is not applicable to Double-Take replication appliances. To update the credentials for replication appliances, use the controller appliance's server properties. See *Controller appliance associations properties* on page 64.

Viewing server details

Highlight a server on the **Manage Servers** page and click **View Server Details** from the toolbar. The **View Server Details** page allows you to view details about that particular server. The server details vary depending on the type of server or appliance you are viewing.

Server name

The name or IP address of the server. If you have specified a reserved IP address, it will be displayed in parenthesis.

Operating system

The server's operating system version

Roles

The role of this server in your Double-Take environment. In some cases, a server can have more than one role.

- **Engine Role**—Source or target server
- **Image Repository Role**—A target for a DR protection job or a source for a DR recovery job
- **Controller Role**—Controller appliance for an agentless vSphere job
- **Replication Appliance Role**—Replication appliance for an agentless vSphere job
- **Reporting Service**—Double-Take Reporting Service server

Status

There are many different **Status** messages that keep you informed of the server activity. Most of the status messages are informational and do not require any administrator interaction. If you see error messages, check the rest of the server details.

Activity

There are many different **Activity** messages that keep you informed of the server activity. Most of the activity messages are informational and do not require any administrator interaction. If you see error messages, check the rest of the server details.

Connected via

The IP address and port the server is using for communcations. You will also see the Double-Take protocol being used to communicate with server. The protocol will be XML web services protocol (for servers running Double-Take version 5.2 or later) or Legacy protocol (for servers running version 5.1 or earlier).

Version

The product version information

Access

The security level granted to the specified user

User name

The user account used to access the server

Licensing quantity

The available and total license count for a Double-Take controller appliance

Licensing

Licensing information for the server

Source jobs

If you are viewing a replication appliance, you will see a list of any jobs from this server. Double-clicking on a job in this list will automatically open the **View Job Details** page.

Target jobs

If you are viewing a replication appliance you will see a list of any jobs to this server. Double-clicking on a job in this list will automatically open the **View Job Details** page.

Associated vCenter/ESXi servers and replication appliances

The vCenter or ESXi servers and the replication appliances that are associated with the controller appliance. To modify the associated appliances, click on the **Edit server properties** link. See *Controller appliance associations properties* on page 64.

Editing server properties

Highlight a server on the **Manage Servers** page and click **View Server Details** from the toolbar. Under **Tasks**, select **Edit server properties**. The **Edit Server Properties** page allows you to view and edit properties for that server. Click on a heading on the **Edit Server Properties** page to expand or collapse a section of properties.

- *Controller appliance associations properties* on page 64—Establishes associations between controller appliances, replication appliances, and vCenter/ESXi servers
- *Server licensing* on page 65—Views, adds, and removes license keys
- *E-mail notification configuration* on page 67—Configures e-mail notification

Controller appliance associations properties

The Controller Appliance Associations properties allow you to manage the replication appliances associated with your controller appliance.

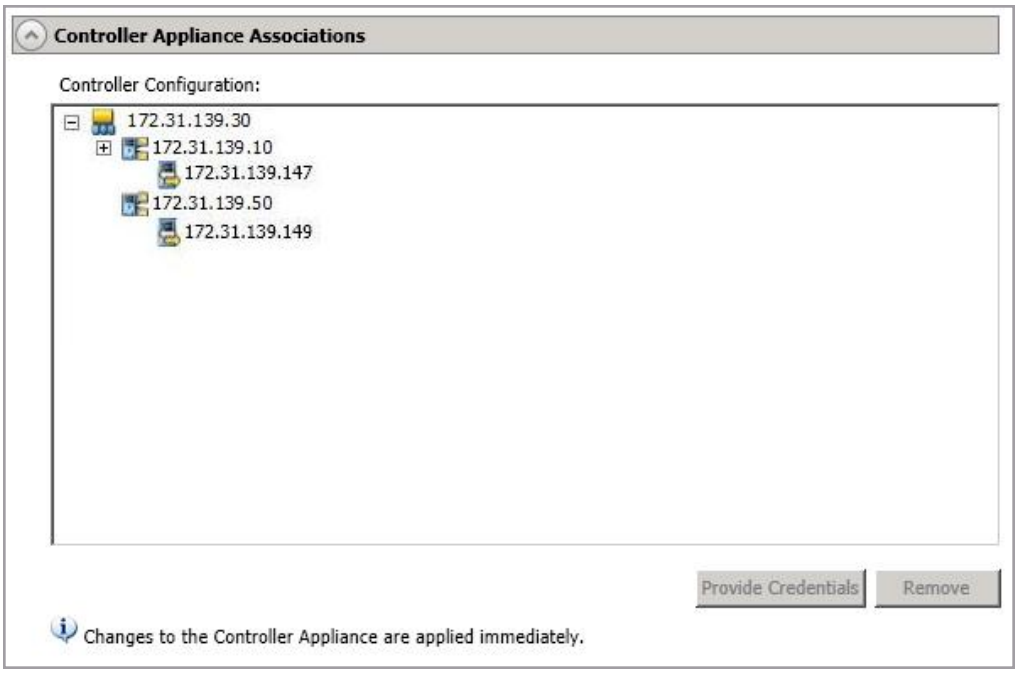

- **Provide credentials**—Select this button to update the credentials for the selected vCenter/ESXi server or the replication appliance. When prompted, specify the **User name**, **Password**, and if necessary the **Domain**.
- **Remove**—Select this button to remove the selected machine. If you remove the vCenter/ESXi server, any associated replication appliances will also be removed. If you remove a replication appliance, just that one appliance will be removed. Keep in mind that you cannot remove appliance associations if a job is using the appliance.

Server licensing

Licensing identifies your Double-Take license keys.

The fields and buttons in the **Licensing** section will vary depending on your Double-Take Console configuration and the type of license keys you are using.

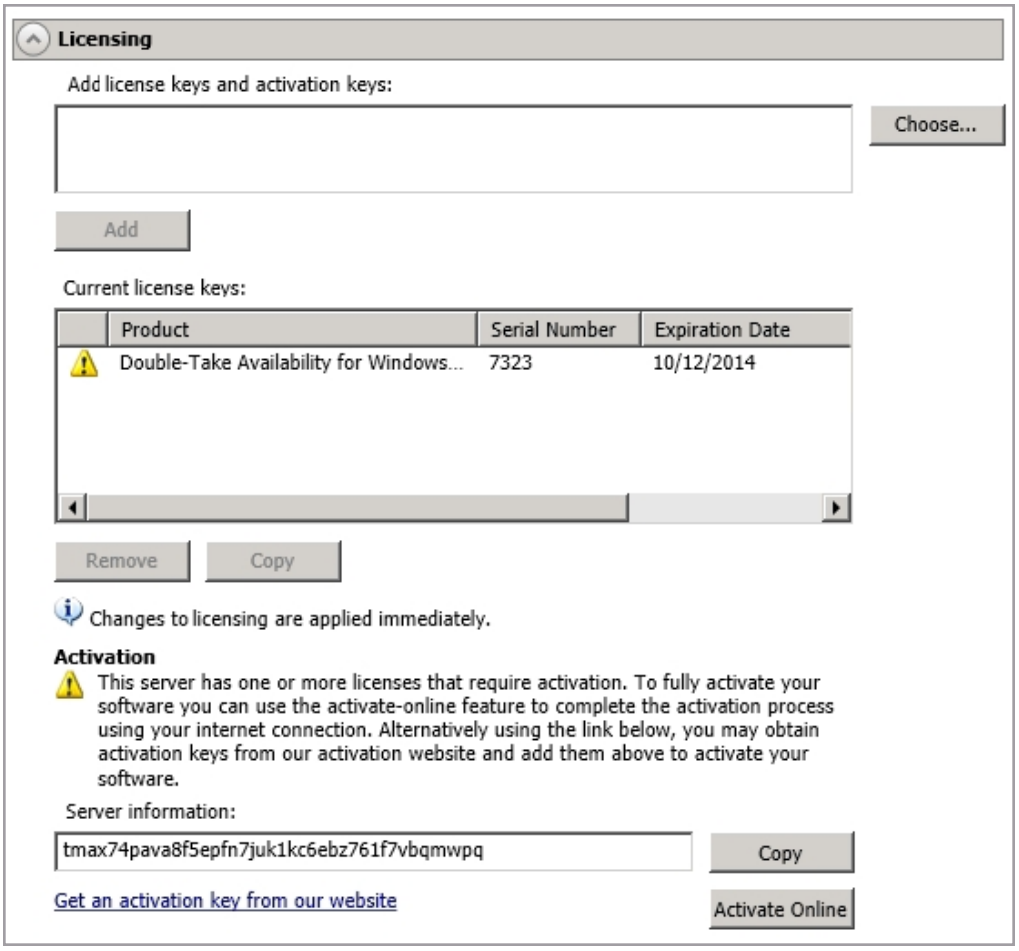

- **Add license keys and activation keys**—Your license key or activation key is a 24 character, alpha-numeric key. You can change your license key without reinstalling, if your license changes. To add a license key or activation key, type in the key or click **Choose from inventory** and select a key from your console's license inventory. Then click **Add**.

  The license inventory feature cannot be enabled if your service provider has restricted access to it.

- **Current license keys**—The server's current license key information is displayed. To remove a

key, highlight it and click **Remove**. To copy a key, highlight it and click **Copy**.

If you are replacing an existing license key that has already been activated, you must remove both the old license key and the old activation key. Then you can add a new license key and activate it successfully. If you are updating an existing license key, do not remove the old license key or old activation key. Add the new license key on top of the existing license key.

- **Activation**—If your activation code needs to be activated, you will see an additional **Activation** section at the bottom of the **Licensing** section. To activate your code, use one of the following procedures.
  - **Activate online**—If you have Internet access, you can activate your license and apply the activated license to the server in one step. Select **Activate Online**. You will not be able to activate a license that has already been activated but has not been deactivated. In that case, you will be prompted to complete a host transfer. Ideally, you should deactivate the license instead of doing a host transfer See the Double-Take Console online help for details on deactivating licenses and host transfers.
  - **Obtain activation key online, then activate**—If you have Internet access, click the hyperlink in the **Activation** section to take you to the web so that you can submit your activation information. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add activation codes and activations keys** field and clicking **Add**.
  - **Obtain activation key offline, then activate**—If you do not have Internet access, go to https://activate.doubletake.com from another machine that has Internet access. Complete and submit the activation form, and you will receive an e-mail with the activation key. Activate your server by entering the activation key in the **Add activation codes and activations keys** field and clicking **Add**.

  The permanent code is specific to this server. It cannot be used on any other server. If the activation code and server do not match, Double-Take will not run.

If your activation codes needs to be activated, you will have 14 days to do so.

If you rename a server that has already has a Double-Take license applied to it, for example if you rebuild a server, you will have to perform a host-transfer to continue using that license. This includes changing the case (capitalization) of the server name (upper or lower case or any combination of case).

- **Total licenses quantity**—If your server is a Double-Take controller appliance, you will see the available and total license count displayed. The controller appliance handles all license management for your agentless vSphere jobs. Each agentless vSphere job will use one license from your available quantity. If you create jobs for five virtual machines at one time, five licenses will be used. When you delete a job, the license associated with the job will be released to be available for another job. Stopping a job will not delete the license associated with the job. If you have no additional licenses available, you cannot create any new jobs. You can add to your available license quantity by purchasing additional licenses.

E-mail notification configuration

You can email Double-Take event messages to specific addresses, using an SMTP mail server. The subject of the email will contain an optional prefix, the event type (information, warning or error), the job name, and the controller appliance name. The body of the message will contain configuration info (date, time, virtual machine names, and so on) and the text of the event messages.

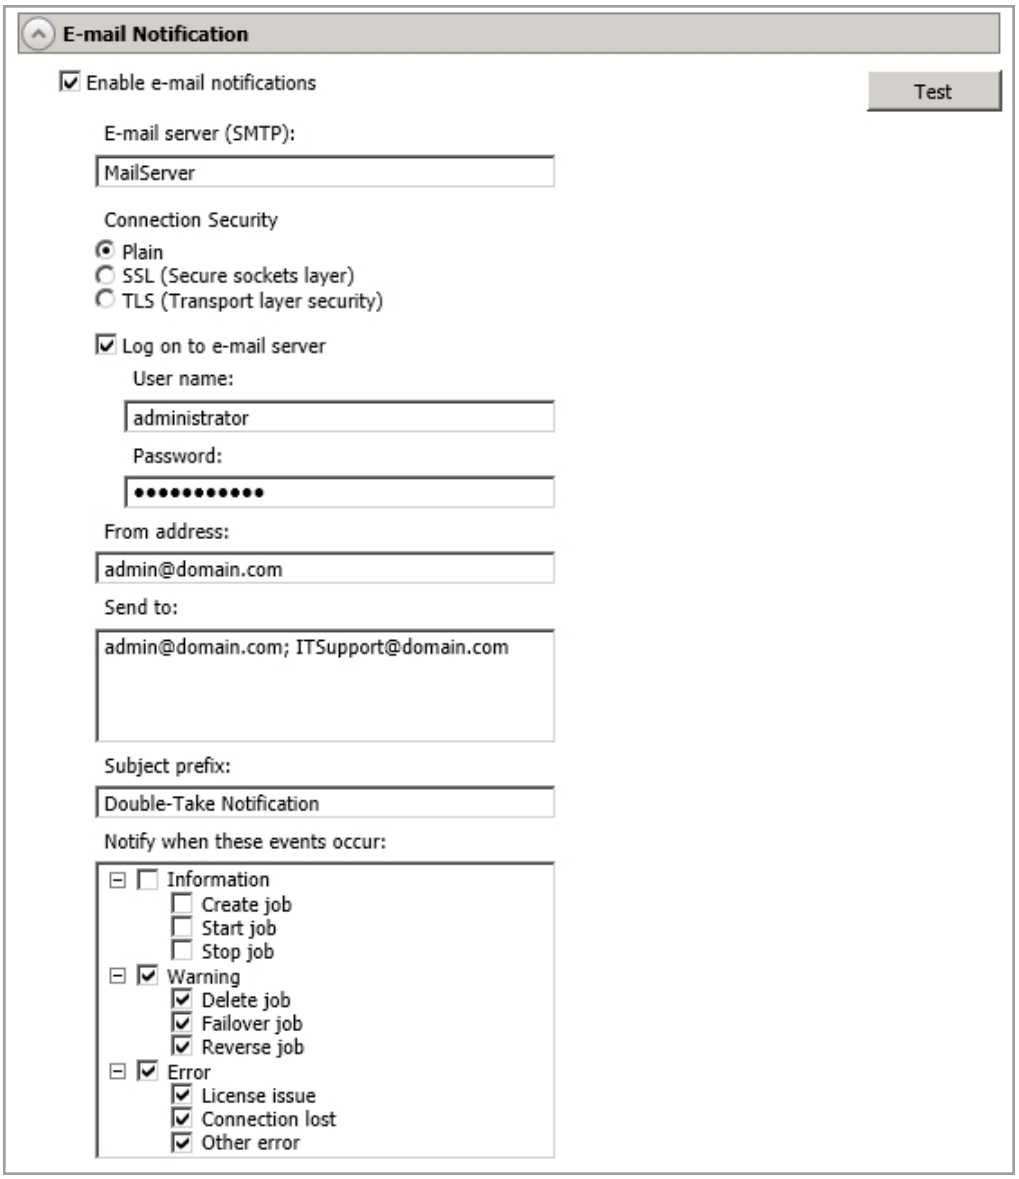

- **Enable e-mail notification**—This option enables the e-mail notification feature. Any specified notification settings will be retained if this option is disabled.
- **E-mail server**—Specify the name of your SMTP mail server. If necessary, after the e-mail server name, you can include a colon followed by the outgoing e-mail requests port number. For example, you might use gmail.com:465.
- **Connection Security**—Select the type of security for your email server. Check with your e-mail service provider if you do not know your security type.

- **Log on to e-mail server**—If your SMTP server requires authentication, enable this option and specify the **User name** and **Password** to be used for authentication.
- **From address**—Specify the e-mail address that you want to appear in the From field of each Double-Take e-mail message. The address is limited to 256 characters.
- **Send to**—Specify the e-mail addresses that each Double-Take e-mail message should be sent to. Enter the addresses as a comma or semicolon separated list. Each address is limited to 256 characters. You can add up to 256 e-mail addresses.
- **Subject prefix**—The subject of each e-mail notification will be in the format Subject Prefix - Event Type, Job=Job_Name, Controller=Controller_Appliance_Name. The subject prefix is optional. The subject line is limited to 255 characters.

  If desired, enter unique text for the **Subject prefix** which will be inserted at the front of the subject line for each Double-Take e-mail message. This will help distinguish Double-Take messages from other messages. This field is optional.
- **Notify when these event occur**—Select each of the events that you want to receive an e-mail notification for. Selecting or deselecting a top level check box (Information, Warning, or Error) will select or deselect all of the items in that category. If desired, you can also select or deselect events individually.

Viewing server logs

You can view the engine and Management Service logs using either of these two methods.

- On the **Manage Servers** page, highlight a server in the list and click **View Server Logs** from the toolbar.
- On the **Manage Jobs** page, right-click a job and select **View Logs**. Select either the source server log or the target server log.

Separate logging windows allow you to continue working in the Double-Take Console while monitoring log messages. You can open multiple logging windows for multiple servers. When the Double-Take Console is closed, all logging windows will automatically close.

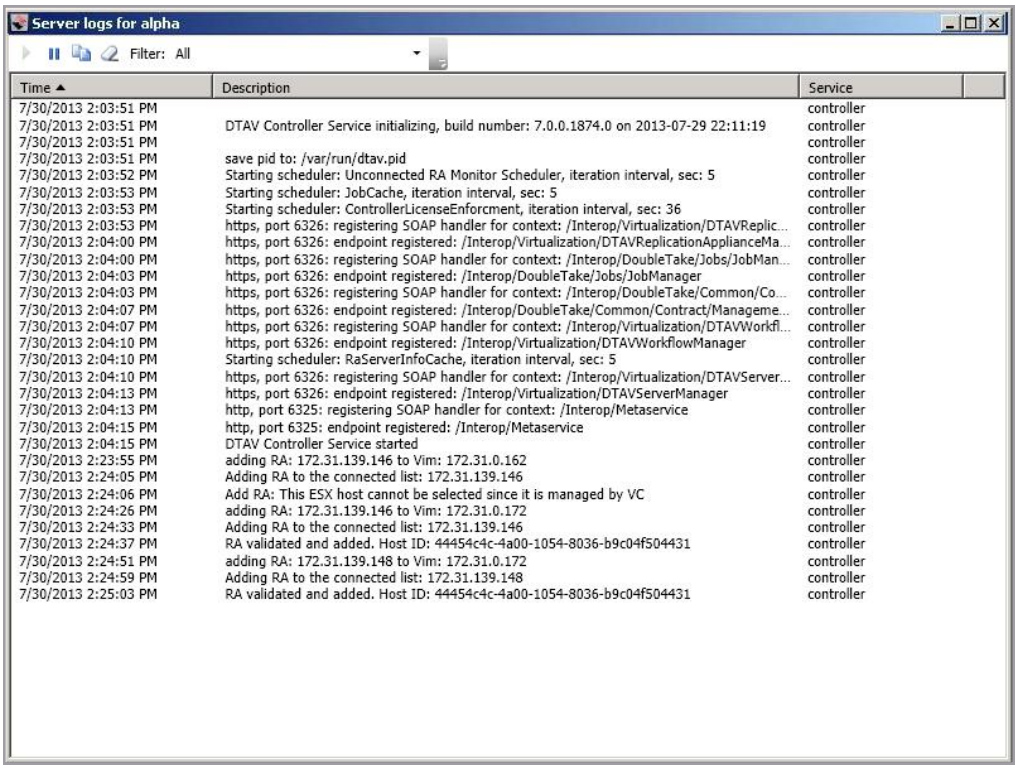

The following table identifies the controls and the table columns in the **Server logs** window.

Start

This button starts the addition and scrolling of new messages in the window.

Pause

This button pauses the addition and scrolling of new messages in the window. This is only for the **Server logs** window. The messages are still logged to their respective files on the server.

Copy 

This button copies the messages selected in the **Server logs** window to the Windows clipboard.

Clear 

This button clears the **Server logs** window. The messages are not cleared from the respective files on the server. If you want to view all of the messages again, close and reopen the **Server logs** window.

Filter

For controller appliances, there is only one log to view so filtering is not applicable.

Time

This column in the table indicates the date and time when the message was logged.

Description

This column in the table displays the actual message that was logged.

Service

This column in the table indicates the message is from the controller log file.